Technology and Innovation, Vol. 20, pp. 117-132, 2018
Printed in the USA. All rights reserved.
Copyright © 2018 National Academy of Inventors.
ISSN 1949-8241 • E-ISSN 1949-825X
http://dx.doi.org/10.21300/20.1-2.2018.117
www.technologyandinnovation.org

PERSON-CENTERED TECHNOLOGIES FOR INDIVIDUALS WITH DISABILITIES: EMPOWERMENT THROUGH ASSISTIVE AND REHABILITATIVE SOLUTIONS

Hemanth Venkateswara, Troy McDaniel, Ramin Tadayon, and Sethuraman Panchanathan

Center for Cognitive Ubiquitous Computing (CUbiC), Arizona State University, Tempe, AZ, USA

Human-Centered Computing (HCC) focuses on a tight coupling of humans in the design of technologies. While HCC has advanced computing research, a new paradigm is needed to better address not only interpersonal variations but also intrapersonal variations that occur over time in order to design more useful and usable solutions for individual users. We propose Person-Centered Computing (PCC) to address users' distinct and ever-changing needs, preferences, expectations, behaviors, and abilities due to aging, geographical location, and contextual factors. There is no greater need for PCC than among individuals with disabilities, where we often find unique requirements that necessitate a person-centric approach. Through co-adaptation, PCC offers a bridge between the development of assistive and rehabilitative applications and design opportunities for the broader population. Our hope is that PCC will motivate accessible and inclusive designs for today's technologies rather than requiring individuals with disabilities to force-fit and modify existing solutions with marginal success. Here we present our PCC paradigm along with two innovative multimedia solutions, the Social Interaction Assistant and Autonomous Training Assistant, to demonstrate how the proposed design strategy not only empowers the target user groups but also impacts the broader population by revealing implicit, unseen needs.

Key words: Person-Centered Computing; Person-centeredness; Assistive technology; Rehabilitative technology; Social Interaction Assistant; Machine learning

INTRODUCTION

As the complexity of computing research advances, technology is becoming increasingly capable of recognizing, adapting to, and accommodating many facets of human behavior. Human-Centered Computing (HCC) is a prominent product of this evolution. HCC has established itself as a field of considerable modern research that focuses on the requirements, tendencies, needs, characteristics, and behaviors of human beings during the research, design, and development of computing solutions (1). For example, solutions that take into account human-centeredness consider and address differences among human cultures and societies within their designs (2). This deeper level of human understanding distinguishes HCC from traditional human-computer interaction practices that treat the user as a static entity and instead focus on general usability. While HCC design has elevated computing research to a new level of compatibility

Accepted: July 15, 2018.
Address correspondence to Troy McDaniel, ASU CIDSE, 699 S Mill Avenue #553, Tempe, AZ 85281, USA. Tel: +1 (480) 694-0021. E-mail: troy.mcdaniel@asu.edu

within differing populations, an even deeper level of understanding is necessitated by the existence of individual variance even within small populations.

Person-Centered Computing (PCC) is a novel paradigm, previously introduced by Panchanathan et al. (3,4), that further enhances the leading principles of HCC (that human populations vary and that technology should adapt to these variations) by applying them at the level of the individual. This strategy dictates that an individual's ideal technological solution is one designed for that particular person, given that he or she is wholly unique. For technology to achieve this level of individualization and, at the same time, be broadly useful and usable to others, PCC hinges upon co-adaptation. Co-adaptation occurs when the user and technology adapt to each other through continual use, with the onus being on the technology to learn and cater to the unique attributes of each individual user.

The paradigm of PCC is inspired by disability research, particularly assistive and rehabilitative applications, where large interpersonal variations in ability are not uncommon among individuals with sensory, cognitive, and/or physical impairments. Large intrapersonal variations over time are also not uncommon. For example, depending on the disease or condition, the impairment may alter with time, changing the way the disability manifests itself. This article presents two examples of technological solutions designed under the PCC paradigm to address the unique needs of individuals with disabilities. The Social Interaction Assistant (SIA) is a real-time vision-based wearable and tabletop social assistive device for recognizing and delivering an interaction partner's non-verbal cues (e.g., facial expressions) to an individual who is blind or visually impaired. The Autonomous Training Assistant (ATA) is a highly-adaptable, individually-aware hardware and software interface capable of delivering multimodal instructions and feedback for at-home rehabilitative training to individuals with varying motor capability and functional goals. The SIA and ATA systems demonstrate that, through deep exploration of the needs of an individual or a few individuals, requirements for design arise that could not otherwise be explored in studies that treat large groups of users as a homogeneous entity; these requirements inspire the creation of more adaptive solutions, as is shown in the examples below.

This paper demonstrates how key findings related to the SIA and ATA support PCC. First, the challenges and perspectives emphasized within the disability space and their implications for modern technology are considered. Next, the research and design process of the SIA are described, as well as some of the key findings and contributions of the project toward improving social interaction for a broader population. The third section highlights the challenges and design underlying the ATA and its novel contributions in motor rehabilitation using a case study application. Finally, thoughts based on the results of these projects and paths for future research in PCC are discussed.

A PARADIGM INSPIRED BY NOVEL PERSPECTIVES FROM DISABILITY

Technology solutions today are still focused on addressing the needs of the 'able' population. The needs of the disabled population are often addressed through modifications and temporary fixes to existing solutions designed for the broader population. For example, although devices like smart phones have been made highly accessible to individuals who are blind, the features of accessibility are built on top of existing platforms rather than the other way around. It is surprising that even though 12.6% (39 million) of the population within the U.S. live with some form of disability (http://disabilitystatistics.org/), technology is still largely designed without them in mind. Economic viability is understood to be the major reason for the lack of innovation in the disability space. To address this issue, there have been efforts over the last decade to make innovation inclusive and accessible. For example, the 'Design for All' paradigm in the context of information and communications technology is a systematic effort to promote universal design in computer and internet-based technologies to avoid standard posterior adaptation procedures that attempt to accommodate the needs of those with disabilities through force fits, workarounds, and afterthoughts (5).

At the Center for Cognitive Ubiquitous Computing, we take inspirations from disabilities to design innovative assistive multimedia solutions where users' needs wholly drive decision-making at every stage (6). Explicit needs are often cognizant and readily available, and therefore straightforward to gather through survey methods such as interviews and questionnaires. Implicit needs are subconscious and represent tacit knowledge, and therefore are more difficult to extract. Our work and experiences in disability research have shown that designing technologies to address the explicit needs of individuals with disabilities paves the way to uncovering addressable implicit needs of the broader population and, through co-adaptation, potentially achieving broader impact and economic viability.

PROTOTYPE STUDY #1: SOCIAL INTERACTION ASSISTANT

Human social interactions are made up of both verbal (speech) and non-verbal (e.g., facial expressions, hand gestures, and body language) cues. A large percentage of interaction (65% or more) is non-verbal (7). Individuals who are blind have limited access to non-verbal communicative cues given their inherent visual nature. The SIA is an assistive aid to help individuals who are blind or visually impaired during dyadic or group interactions by providing access to non-verbal information. The SIA consists of three components: sensing, processing, and delivery. Sensing involves capturing information of the interacting partner(s) and their surroundings through sensors (e.g., visual and audio) embedded on the user and in the environment. The processing component mines the data gathered from the sensors for patterns of non-verbal information. The interaction phase discreetly delivers the processed information in the form of cues that are relevant and useful to the user. Rather than interfere with an ongoing social interaction, these cues augment the user's awareness of their interaction partner.

Role of Non-verbal Cues in Social Interaction

Social interactions play an important role in helping us communicate effectively. They form an integral part of our everyday communication that enables us to convey our emotions and feelings in an efficient manner. Besides facilitating communication, they also help us socially connect with friends, family, and peers. It is well understood that such social connectedness aids in reducing stress by providing emotional and psychological needs, which has major health benefits. Hence, social interaction is essential for a higher quality of living. Given the visual nature of non-verbal social cues, individuals who are blind are at a disadvantage when interacting with their sighted peers and colleagues. It is not uncommon for this population to feel awkward or embarrassed in certain social situations due to miscommunications that could have been avoided if they had had access to visual information. Such situations can lead to social avoidance and, eventually, social isolation, which can be detrimental to a person's productivity, health, and wellbeing (8).

Components of the SIA

The SIA, depicted in Figure 1, consists of sunglasses with a discreetly embedded camera (a tabletop web cam may be used for social settings at a table); a computing device (i.e., mobile phone) to process the images captured by the camera; and haptic devices for information delivery, including a vibrotactile belt when user is mobile as well as the haptic face display when user is stationary and seated at a table. The haptic face display (9) consists of a two-dimensional array of vibrotactile actuators that are driven by the processing device. In Figure 1, the haptic face display has the form factor of an ergonomic mesh chair. The user seated in the chair experiences vibrotactile stimulation on his or her back, and it is the dimensions of the vibrations that encode the non-verbal cues. We are currently exploring other form factors for two-dimensional vibrotactile displays, including a wearable vest for mobile use as well as a tabletop display for active exploration. The chair's vibrotactile display consists of 3.3 volt DC eccentric rotating mass (ERM) pancake motors arranged in the form of an array of six rows and eight columns. The ERM motors are spaced 2 cm apart horizontally (center-to-center) and 4 cm apart vertically (center-to-center). The SIA also has a haptic waist belt (10) to provide cues regarding position and proximity of interaction partners to augment situational awareness in mobile contexts. The belt consists of seven ERM motors placed equidistantly along the length of the belt, ending at the left and right sides of the torso, with the middle vibration motor aligned with the midline of the body. The motors vibrate discreetly to indicate the relative orientation of the interaction partner using body site and interpersonal distance (proximity) using vibration duration or rhythm. These cues assist users in orienting themselves to face their interaction partners and maintain the appropriate distance from them. To allow this system design to serve the needs of each individual, both the complexity of tactile feedback and the capture of facial features are adapted using swappable mapping formats and domain adaptation, respectively, as described below.

Adaptation of Information Delivery for Person-centeredness

Discussions with individuals who are blind and visually impaired revealed a desire to virtually explore the faces of their interaction partners to experience and understand facial expressions first hand. Other users were merely content with an assistive aid that recognizes and conveys to them the emotions of their interaction partners. It was clear that each user's needs were unique, and a one-size-fits-all policy was not going to be universally acceptable to the blind community. It was necessary to build a system that could adapt to the needs of each individual user. Since the delivery component of the SIA is the interface between the user and the system, it also has the maximum scope for adaptation. The delivery phase of the SIA incorporates person-centeredness by providing the user with output cues of varying granularities. We outline three adaptive delivery mechanisms for the SIA that enhance person-centeredness.

Literal Mapping

In this design, there is a direct mapping between the 49 landmark features displayed in Figure 2 and the vibrotactile actuators (regions of the haptic face display). In this setting, users passively explore the faces of their interaction partners in a virtual manner and gather information about the facial expressions first hand. The vibration motors in a region are actuated based on the change in the position of key landmarks relative to other landmarks. Tracking the relative change helps to account for overall movements of the face, which, in turn, changes the position of all

Figure 1. Social Interaction Assistant components and technologies.

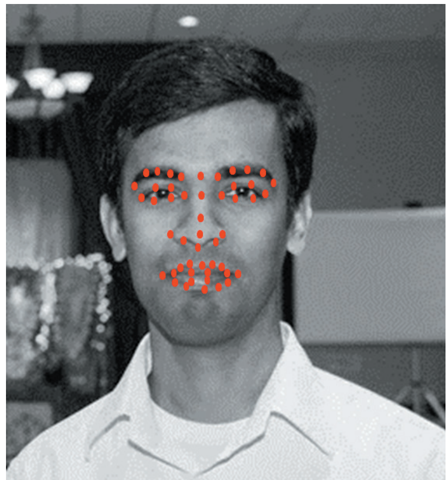

Figure 2. Facial landmarks for literal mapping of features.

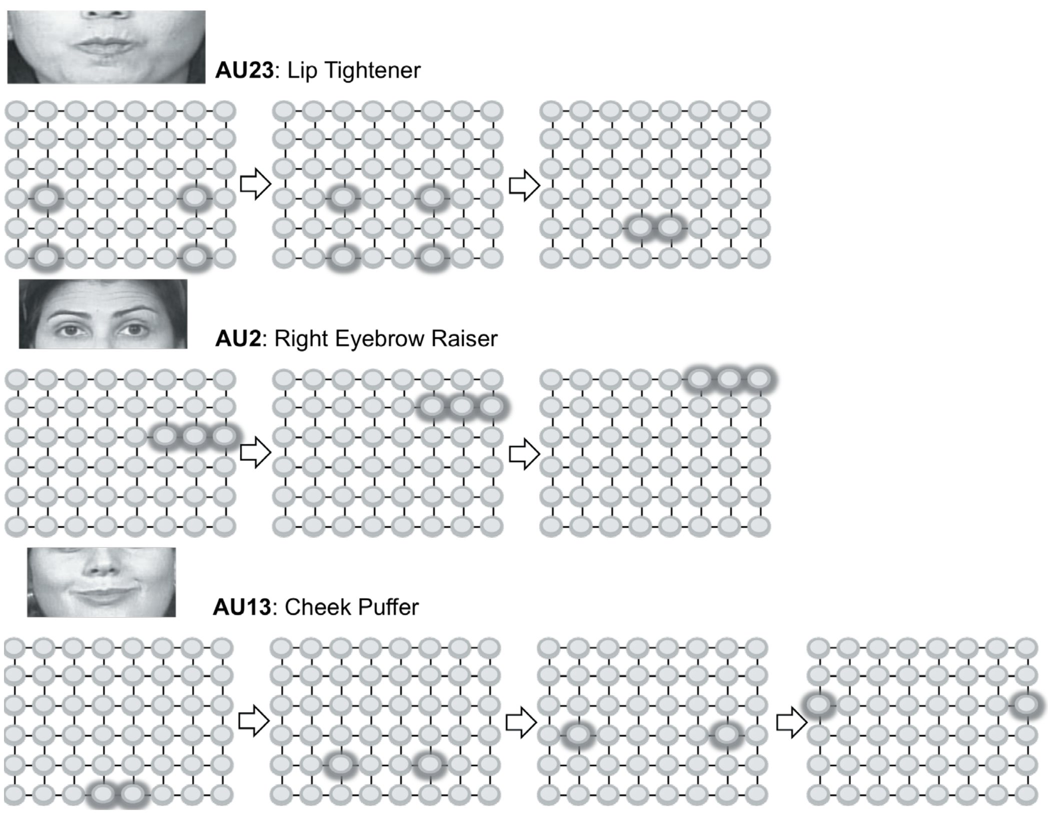

Figure 3. Semi-literal mapping of facial action units to vibrotactile patterns. The array of dots below the action unit represents the six by eight two-dimensional vibrotactile display of the haptic face display. The darkened dots indicate actuated vibration motors, and their temporal sequence is from left to right, which encode the action unit represented by the face images. Three example patterns from a larger set are shown.

the landmarks. The landmark movements arising from external factors such as camera movement and head movement need to be accounted for to capture only relative landmark movement. A minimum set of landmarks is chosen to encode facial expressions. This is the most direct mapping of facial expressions into vibrotactile patterns, and the onus of recognition, interpretation, and decision making is on the user.

Semi-literal Mapping

In this mapping, a layer of abstraction is introduced, compared to the direct mapping of facial landmarks, by using facial action units. The Facial Action Unit Coding System is a categorization of facial muscle movements commonly used by psychologists to study human facial expressions (11). In earlier work (12), we designed and evaluated vibrotactile spatio-temporal patterns to intuitively communicate a subset of facial action units that are most prominently featured in the six basic emotions (13): happiness, sadness, surprise, anger, fear, and disgust. The haptic face display was divided into three regions: upper region to transmit eyebrow movements; middle region to convey movements of the nose and cheek; and lower region to map the movements of the lips. From an initial set of vibro-tactile patterns, pilot testing allowed the selection of a subset of distinct and intuitive mappings, which were then further evaluated to more deeply explore their naturalness and distinctness. This study yielded recognition accuracies in the range of 80% to 96% (with a trial average 91%). It was estimated that the vibration pulse length, ranging from 250 ms to 1,000 ms, had little effect on recognition accuracy. With no significant difference found between different pulse durations, it was concluded that the perception of the proposed spatio-temporal patterns was not affected by the variation in individual pulse widths. As a result, pulses of short duration could be used to enable higher throughput and improved communication rates. Figure 3 depicts a few of these action units and the corresponding vibration patterns. During real-time presentation, when multiple action units are detected, the haptic face display will convey these sequentially or, alternatively, will convey the most confident (intense) action unit.

Symbolic Mapping

This mapping provides the highest level of abstracted information to the user in the form of emotions. Whereas in the previous mappings, the user made the final decision about the emotion of their interaction partner, in this representation, the SIA processes the sensory information from the camera to determine the basic emotion of the partner and subsequently conveys a symbolic representation of the emotion to the user. This mode of operation can be useful when the user is overloaded with sensory information from other sources and is merely interested in knowing the emotional state of the user without having to devote any cognitive processing to gather and analyze such information.

To select effective and efficient patterns to convey emotional states, we explored a large design space of spatio-temporal vibrotactile stimulations that could potentially elicit emotions of happiness, sadness, surprise, anger, fear, and disgust in the user (9). Beginning with a design space of 150 patterns, pilot testing revealed 54 patterns as distinct and potentially useful. These remaining patterns, which also included timing variations (various pulse widths) and saltation versions (14) of selected patterns, were evaluated for their effectiveness at evoking emotions. A study of consensus among participants yielded 20% to 30% agreement for certain patterns. The results were averaged across gender and vibration duration. This yielded a consensus of 26% for anger when using the Six Motor Burst pattern; 28.1% for happiness when using the Snake pattern; and 22.9% for neutrality when using the Spiral pattern—that is, the absence of emotion, which may be useful to convey between long periods absent of emotional content (Figure 4). We hypothesize that the Snake pattern was interpreted as playful because it feels like someone running a finger across the user's back in a winding pattern and that the Six Motor Burst pattern angered participants due to its intensity and randomness. Moreover, it was observed that patterns that lasted longer elicited sadness, and patterns that were shorter were more likely to evoke happier emotions.

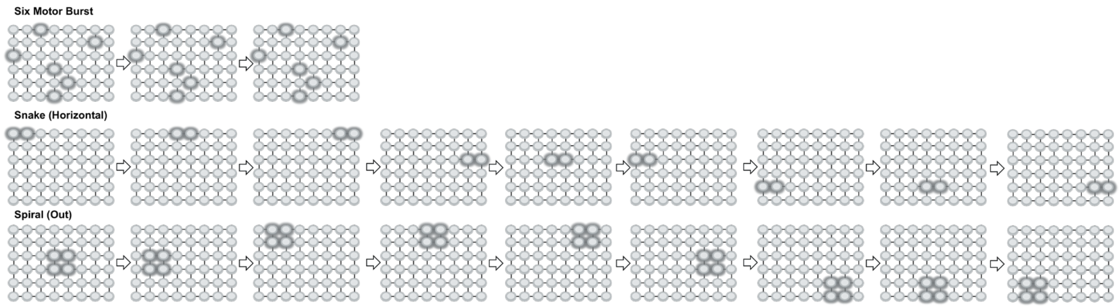

Figure 4. Symbolic mapping of three emotions to vibrotactile patterns. This approach aims to effectively and efficiently convey an interaction partner's emotions by evoking those same emotions in the user. This strategy is to enhance communication by enabling the user and his or her partner to share non-verbal emotional experiences.

Adaptation of the Algorithms for Person-centeredness

Based on a user's needs, they may switch between the different levels of abstraction (literal, semi-literal, or symbolic) of the SIA. The delivery phase of the SIA receives its inputs from the processing phase of the SIA. At the core of the SIA is the processing algorithm that analyzes the data captured through the sensors. The processing algorithm outputs landmarks, action units, and emotions of the interaction partner, which are conveyed by the delivery phase to the user on an as-needed basis. The processing phase consists of computer vision algorithms for face detection, facial landmark extraction, action unit recognition, and facial expression recognition. These algorithms analyze the face of the partner to extract this information in real time. The computer vision algorithms are trained to make these predictions using a large pool of labeled data using supervised learning—a paradigm of machine learning. Algorithms trained in this manner can then be applied to make predictions on test data consisting of previously unseen faces.

Supervised learning algorithms often make the implicit assumption that the test data is drawn from the same distribution as the training data. Such algorithms become ineffective when these assumptions regarding the data are violated. For example, a facial expression recognition algorithm trained on white facial images is likely to degrade in performance when tested on facial images from a different ethnic group, say African American. This is because of the change in data distributions of the train and test data—the facial features are different, the facial expressions could be different based on cultural dissimilarities, etc. In addition, the paucity and/or poor quality of data from a distribution also limits the efficacy of the classifier. When the number of samples or their quality is limited, supervised learning algorithms fall short in learning a well-generalized classifier. These kinds of problems are addressed using transfer learning and domain adaptation techniques. Transfer learning and domain adaptation algorithms extract knowledge from one or more tasks or domains and utilize (transfer) that knowledge to design a solution for a new task or domain. These adaptation procedures transfer knowledge from a source domain (distribution) to a target domain (distribution), in the form of learned models and efficient feature representations, to train well-generalized classifiers for the target domain (15,16).

In designing the SIA, we have applied domain adaptation to make the processing of action unit and facial expression recognition robust across a wide range of face types. Facial image data can vary due to the following reasons: differences in image resolution, brightness, occlusion, changes in camera point-of-view, and inherent diversity of the sample space. It is impractical to train individual recognition models for every kind of distribution due to a paucity in labeled data, the cost involved in training, and the wide variety of distributions that arise from vision-based data. Therefore, domain adaptation-based solutions are applied to adapt models trained on one domain (distribution) to other domains (distributions). In the SIA project, we have explored domain adaptation at the classification stage with the use of linear models.

In the Coupled Support Vector Machine (C-SVM) algorithm (17), we train a pair of SVMs–one for the source domain and another for the target domain–with limited labeled data in the target domain. In the Nonlinear Embedding algorithm (18), we learn to adapt the two domains by projecting the features into a common subspace using a nonlinear embedding model. We test the facial expression recognition algorithms using standard datasets, such as CKPlus (19) and MMI (20), and demonstrate how domain adaptation improves the accuracies of prediction on the target domain when there is very limited or no labeled data. The algorithms for the processing unit of the SIA are trained on standard existing datasets, such as CKPlus and MMI, but will be made robust with domain adaptation to recognize the facial expressions of a wide range of faces and adapt to the social group of the user.

PROTOTYPE STUDY #2: AUTONOMOUS TRAINING ASSISTANT

Need for Person-centeredness in Rehabilitative Motor Learning

Motor rehabilitation is perhaps one of the most well-aligned domains for person-centric computing. In this field of research, new devices, systems, and techniques for the enhancement of rehabilitative exercise are in constant development. This is no surprise, as rehabilitative care is one of the most expensive fields of medical expenditure in the U.S. every year (21). Technologies researched in this field are often intended to assist the user with rehabilitative exercise by assessing performance, providing feedback, and creating an interactive space more conducive to motor training than interacting with ordinary equipment in a household or clinical environment. This technology may also be responsible for detecting and reacting to changes in user state, including fatigue or compensatory behavior while exercising. Results reported related to the usage of this technology often include an increased sense of autonomy, improved motivation to exercise, higher levels of dedication to regular exercise, and, by consequence, improved health outcomes.

However, most of the technologies designed for these purposes often fall short of achieving significantly positive outcomes outside of a very narrow group of test users and subjects. A significant cause of this shortcoming is that an overwhelming majority of research efforts in the field treat users of rehabilitative technology as a homogeneous entity. Quite to the contrary, rehabilitation is a field brimming with individual variance. Each individual case involves a variety of experiential, physical, and human factors, including an individual's motor ability, level of impairment, location of impairment, functional goals, muscle strength, progression through therapy, age, gender, motor exercise regimen, and training environment, among many others, all of which directly affect design choices necessary to create an effective solution for that individual (22). Each individual, based on these factors, should be considered a unique entity.

Therefore, it is argued that research can more effectively be conducted in rehabilitation when it begins with a single person and a real challenge posed by that person. One of the most immediately apparent flaws with this approach is that, by focusing on a single individual, an effective solution may be created for that individual, but it would be too specialized to be useful for other individuals with significantly different challenges. Fortunately, using the principle of co-adaptation, technologies that begin with a single individual can mold themselves to meet the varying needs of other individuals, as is demonstrated in the design process of the ATA system, which is detailed below.

Supervised vs. Unsupervised Motor Learning and the Role of the Physical Trainer

As the range of application areas for rehabilitative technology is rather broad, it is important to establish the scope of the work presented here. Two forms of rehabilitative training are considered for this purpose: supervised and unsupervised training. In supervised motor training, a physical trainer, therapist, or other expert is present to perform some (or all) of the following functions during an individual's training (23): assess the individual's performance at a motor task; provide feedback to the individual on his or her performance, clearly distinguishing between what is done well and what needs improvement; adjust characteristics of the motor task so that it matches

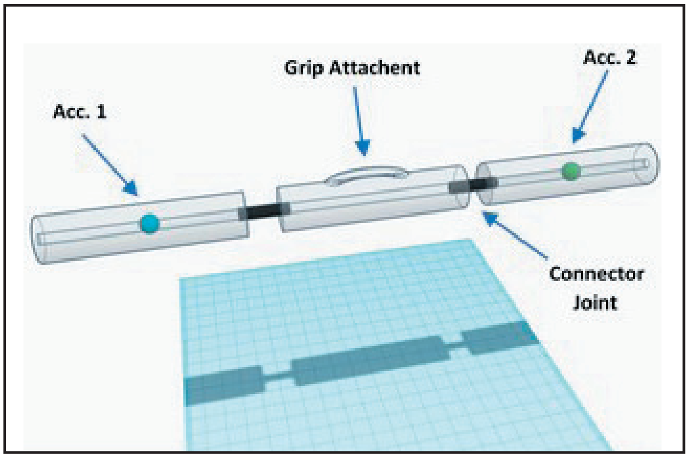

Figure 5. Intelligent Stick early prototype design sketch.

the individual's current level of skill; maintain and update goals for performance on the task; respond to an individual's emotional status, including interest in the task, boredom, or frustration; correct unsafe behavior such as overexertion or compensatory motion; react to potentially dangerous changes in the individual's physical state, including fatigue or dizziness; and encourage and motivate the subject.

While these tasks are common among physical trainers during supervision of rehabilitative exercise, there is no universally accepted standard for the way in which these practices are implemented. As such, trainers use a myriad of techniques, across a variety of interaction channels, to deliver interactive, meaningful rehabilitative exercise to subjects. In addition, the amount of weekly exercise required for an individual to maintain steady progress in a rehabilitative program typically exceeds the availability of a trainer or therapist, which leads to the assignment of unsupervised training (often in the individual's home environment). In this form of training, an individual is typically expected to complete a set of rehabilitative exercises without expert supervision. Accordingly, the above services that an expert would provide in supervised training are missing in many of the traditional forms of unsupervised training, wherein the subject either uses simple training equipment or exercises without any equipment or interface. In unsupervised training, an automated solution for the delivery of supervisory tasks is therefore highly desirable and is the focus of this work.

Introduction and Overview of the ATA

The ATA is an interactive system for guided motor training in the home environment. It consists of the following components:

- The Intelligent Stick, a rod-shaped training device equipped with various components for sensing and feedback, including vibrotactile motors, an accelerometer, and a gyroscope. The design of the device is modular, allowing for various parts to be swapped in and out freely. This design strategy follows Universal Design principles, allowing the device to adapt itself to the needs of users with a variety of grip strength, form factor, exercise task, and feedback requirements. A design sketch of this device is featured in Figure 5.
- A Microsoft Kinect sensor for real-time skeletal joint tracking and facial tracking for affective response detection. Motion data from this sensor is fused with data from the Intelligent Stick accelerometer and gyroscope in real time to track and assess a user's motion during exercise, thus alleviating the occlusion problem posed by the static camera sensor.
- An audiovisual training interface designed as a serious game using stealth assessment. The

type of game as well as the design depends on the interests and motion tasks of the subject. An example implementation of a game (Island Fruit) is covered in previous work by Tadayon (24), wherein the game adapts not only to the subject's motor performance but also to the subject's emotional response (affect) seamlessly in real time.

To facilitate at-home training, the system requires the participation of a subject and trainer. The system is deployed in the subject's home environment and tracks the performance of the subject during use, which it reports automatically in regular intervals to the subject's trainer. Based on this report, the trainer then assigns new exercises and performance goals to the subject via a remote interface, which are then uploaded back to the system through cloud storage. The system updates itself and implements the latest motor tasks and goals on the next use. Adaptation for person-centeredness in the ATA is implemented in three phases: the creation of a flexible motor assessment model that can account for trainer variability; the implementation of multimodality through the audio, visual, and haptic channel in feedback; and the interweaving of a variety of game archetypes with an individual's specific training protocol using stealth assessment.

Case Study Design, Individual Attributes of Learner

The design of the ATA study stemmed from a case study involving a single subject. The subject is hemiparetic, having full usage of one arm with partial physical impairment in the other. The subject's physical trainer utilizes martial arts training as a context for rehabilitative therapy. The challenge was to develop a system that could allow the subject to receive guided martial arts training at home using the trainer's expertise and training regimen. In other words, the system would need to provide the services of supervised training highlighted above while also meeting the following requirements specific to the subject: The subject should use the nonparetic arm to guide the paretic arm; the subject should be able to easily use the system without complex or difficult setup; the subject should be allowed to freely move the arms during the exercise; the subject's equipment should include a mechanism to secure the subject's grip; and the interface should include some form of gameplay related to the subject's martial arts exercise to help motivate the subject. Our approach in this case consisted of several phases, which are described next.

Phase 1: Assessment

In this phase, supervised training sessions between the subject and trainer were observed to determine the training protocol and trainer's method for assessment and feedback delivery. This process is detailed in prior work by Tadayon et al. (25). Based on these observations, it was determined that the subject would be assessed on motor performance using three metrics: posture (body positioning), pacing (rate of motion), and progression (degree of motion). The targeted values for each of these categories were provided by the trainer, as well as a series of motion exercises ranging from simple tasks (elbow flexion/extension) to more complex tasks (swing motions using stick equipment). To capture these motions in real time at home, the Intelligent Stick device described above was created. The stick form factor was chosen, as it matched the form factor of the stick equipment the subject was already using in live training, but the design was modularized to adapt to other individuals and training regimens. The accelerometer and gyroscope on the Intelligent Stick track and report the user's motion (progression) and rate of motion (pacing) to the system in real time, while the added Kinect sensor tracks joint motion across the body (posture). Combined, these two devices are used to assess the subject on the same criteria used by the individual's trainer.

The three-category framework for assessment used in the case study (posture, progression, and pacing) captures both the temporal and spatial properties of an individual's motor performance during training and was derived from the observation of a single trainer's protocol. Similarly, the Intelligent Stick design was inspired by the equipment used by the case study subject and included a wrist strap mechanism to ensure that the subject's paretic arm remains fastened to the device despite weak grip strength. In both design and implementation, this

Figure 6. Island fruit game interface.

system was entirely developed from the observation and interaction with these two individuals, yet it is evident that the framework and technology can adapt to various other cases of motor training. Swappable components ensure that the Intelligent Stick's length, width, onboard sensing features, onboard feedback features, and other details can be adapted to match other cases of rehabilitative training. The components of the assessment framework described above can map to various assessment standards commonly used in practice, such as the Wolf Motor Function Test (26) or the Barthel Index (27). Hence, person-centric research need not restrict the applicability of a solution across the breadth of a problem space.

Phase 2: Feedback

In the second phase, feedback delivery mechanisms were incorporated so that the system could deliver real-time guidance to the subject during exercise. Three modalities of feedback were observed during live training between the subject and trainer: audio, visual, and haptic. Audio feedback was delivered as verbal instruction from the trainer, while visual feedback was given by demonstrating the correct postural form to the subject. Haptic feedback consisted primarily of the use of "push" and "pull" cues from the trainer to guide the subject's joints to the correct positioning during motion. The system's on-screen game interface is capable of providing audio-visual feedback, and haptic feedback was incorporated via the addition of a vibrotactile response module to the intelligent stick. This module includes several onboard vibrotactile motors, which were added to the stick in the areas gripped by the subject during motion. These motors could deliver haptic cues to imitate the guiding hand of the trainer when the system deems it necessary during at-home training. The specific patterns delivered across this vibrotactile band were determined by the trainer to ensure that they were consistent with the individual's training.

Phase 3: Game Design

In the final phase of the project, the research team worked closely with the subject and trainer to determine a game design that would best fit both the interests of the subject and the motion tasks. It is well known that there is no one-size-fits-all approach to the design of a serious game; just as individuals have varying tastes in art, music, or other mediums, interests in game type vary greatly. Hence, the selection of a best-fit game type and gameplay elements is very much a person-centric process. In this case, factors that contributed to the choice of game included the subject's preferences (determined via interview), age, and motor ability, as well as the specific motor tasks chosen by the trainer (stick swing motions). Based on these criteria, an Island Fruit game depicted in

Figure 6 was chosen, wherein the subject slices fruit using a virtual sword representation of the Intelligent Stick. In general, to utilize a person-centric approach, the design of the game interface consists of the following phases requiring the participation of both the subject and trainer: a selection of motor tasks to be performed, a selection of metrics that serve as evidence for the completion of those tasks, a selection of game archetype that fits the level of complexity of the task that needs to be performed, and a mapping from the motor task performance evidence to game evidence within gameplay objectives.

Once the game's context was chosen, the next step was to incorporate the assessment and feedback frameworks above into its design. A critical requirement here is that the system's assessment and feedback should not interfere or distract the subject from gameplay, as this would increase cognitive load and deter player engagement (28). To accomplish this, motor feedback is interwoven with gameplay in an invisible manner using Shute's method of stealth assessment (29). Fruit objects in the game represent critical points along a motion trajectory that a subject must contact to satisfy the trainer's requirement for progression. The subject must hold both hands on the stick during the game so that the nonparetic arm can guide the paretic arm in a bimanual swing motion; if this is not the case, the in-game virtual sword wobbles in place and cannot be used, satisfying the requirement for posture. Finally, the fruit objects move through the air at a speed dictated by the trainer's requirement in pacing, and, to strike the fruit, the subject must move at the correct pace. Haptic feedback from the Intelligent Stick is given when the subject contacts the center of a fruit object, indicating an optimal swing trajectory. Visual observation of the virtual sword's movement through the screen allows the subject to form the correct posture. A sound effect on contact between fruit and sword indicates that the user's pacing is accurate. Using the subject's contact with fruit objects, as well as the data from the Kinect and Intelligent Stick, the system can seamlessly assess and provide feedback to the user without distracting from gameplay.

Furthermore, adaptation is implemented to maintain a level of optimal challenge for subject engagement. Dynamic Difficulty Adaptation is critical in serious games because it helps account for skill variability both within a single-subject over time, as in the case study here, and between different subjects, if applicable. In this case, using metrics provided by the trainer, the system maintains a performance profile of the subject, which is updated in real-time. Based on this performance measure, as well as the subject's emotional response to the gameplay (estimated using facial emotion tracking from the Kinect's video feed), the system adapts game difficulty by means of adjustment of the size and speed of fruit objects.

Studies on Assessment, Feedback, and Stealth Assessment/Adaptation

Throughout each phase of the case study, evaluations were conducted to determine how well the ATA system matched the requirements of the subject and trainer in that phase. The results of these evaluations directly resulted in design decisions at each phase to improve the system's fit to the individual's motor ability and the trainer's exercise program. Two of these evaluations are summarized as follows.

Multimodal Feedback Evaluation

For this evaluation, detailed in Tadayon's work (24), the goal was to determine how well the subject was able to perform a two-minute task under three different multimodal feedback environments. After the optimum form for a motor task (umbrella motion) was demonstrated by the trainer, the goal of the subject was to consistently maintain this optimal trajectory while repeating the motor task using the Intelligent Stick device. In the control condition, no feedback was given by the system, and the subject simply swung the Intelligent Stick repeatedly in an umbrella pattern. In the "Sigrist" condition, the subject received haptic guidance from the stick, audio feedback on pacing errors, and visual feedback on posture, as inspired by a review from Sigrist et al. (30). In the final condition, the subject was allowed to choose how feedback was given by the system. In this case, the subject chose to receive only haptic feedback on motion trajectory (progression). The subject's error in each session was determined as the deviation of the centroid of the stick from the expected location in space-time, estimated in accelerometric

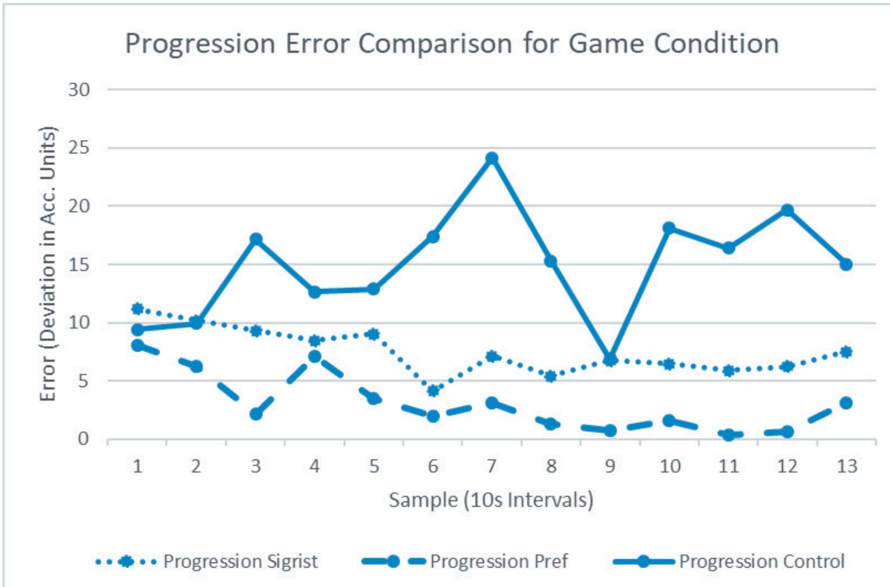

Figure 7. Progression error over 2-minute umbrella task.

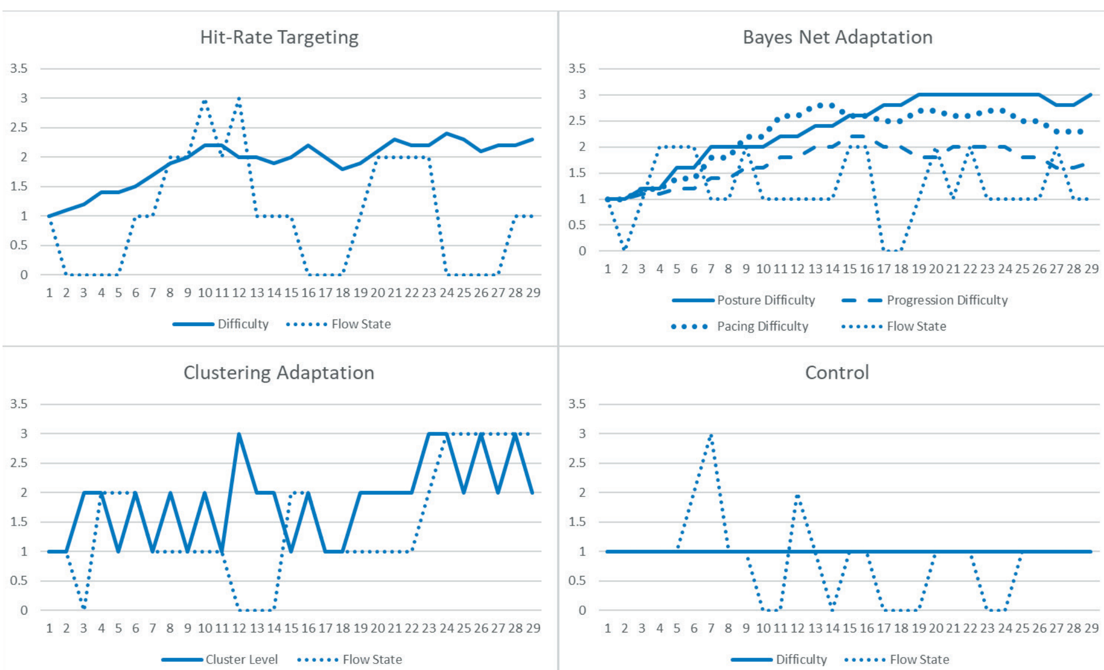

Figure 8. Flow state progression for case study subject. For flow state, a value of 0 corresponds to unknown state, while 1, 2, and 3 correspond to boredom, flow, and anxiety, respectively. For difficulty, the value is a normalized representation of the level of difficulty presented to the user during each time sample.

units. Interestingly, the results, as shown in Figure 7, indicate that the highest reduction in error, or the greatest improvement in performance over the two-minute session, occurred for progression when the subject chose the feedback to be received. While these results may not necessarily transfer to other individuals, they indicate that individual preference may play a role in the optimal mapping of modalities to feedback in a multimodal exercise environment, a notion that merits future evaluation among a variety of subjects.

Stealth Adaptation Evaluation

In this evaluation, also detailed in Tadayon's previous work (24), several different adaptation strategies were implemented for the Island Fruit game described above. The goal was to determine which adaptation strategy yielded the best affective response from the subject. The first approach, hit-rate targeting, attempted to maintain a certain number of fruits sliced by the user for each deployment in the game. Out of three fruit objects, for example, the system tuned the difficulty to ensure that the user could slice two. If the user sliced less than two, difficulty was lowered, while if the user sliced more, difficulty was increased. The clustering approach observed the user's performance in multiple categories (posture, progression, and pacing), and grouped the user into a mastery level or cluster based on this performance. The system then adapted the difficulty to match the user's level of mastery based on the assigned cluster. The Bayes Net adaptation strategy used a Bayes net to maintain belief states about the user's performance independently in each of the three performance categories and adapted the game for each category independently (for example, if the user was only lacking in pacing performance, only the speed of the fruit objects was lowered while other parameters remained the same). Finally, the control condition incorporated no adaptation strategies.

Using the Kinect camera and Visage library, an estimate of the user's facial expression was captured at each ten-second interval. Facial expressions corresponding to the state of flow (31), anxiety/frustration, boredom, and unknown expressions, are captured in Figure 8. Flow-state ratio corresponds to the portion of the time that the subject was considered to be in flow. Each adaptation strategy was incorporated in a five-minute session during which the user played the Island Fruit game. As shown in Figure 8, the highest flow-state ratio was yielded by Bayes Net adaptation at 0.3, followed by hit-rate stabilization at 0.233, clustering at 0.2, and control at 0.067. A separate study by Tadayon (24) indicated that the Bayes Net approach also yielded the highest overall player performance, validating the notion that higher engagement via flow leads to better learning (32).

CONCLUSIONS

Our interaction with individuals with disabilities has inspired a person-centric computing paradigm to address the unique needs of this target population within the context of assistive and rehabilitative applications. We have found that through targeting the explicit needs of this community, solutions that address the implicit needs of the broader population may be advanced; for example, here we presented two person-centric multimedia solutions, the Social Interaction Assistant and Autonomous Training Assistant, which could potentially augment remote communications across long distances and exercise or sport applications, respectively, for the general population. While the terms abilities and disabilities have been used throughout this paper, it is our view that ability is a continuum, ranging from disability to ability to super-ability. In other words, we are all disabled to some extent and fall somewhere on the aforementioned continuum. Therefore, person-centeredness not only facilities the design of more useful and usable technologies for disabilities but also empowers individuals throughout all degrees of ability. The case study approach utilized here reflects the idea that person-centric research must begin with the needs of the individual while adapting itself to the needs of the many. These individual needs are highlighted as well as an evaluation that serves as a proof of concept. Future work in this space will apply the approach toward a larger variety of individuals to indicate how the adaptation mechanisms present in the systems can perform as these needs change.

ACKNOWLEDGMENTS

The authors thank Arizona State University and National Science Foundation for their funding support. This material is partially based upon work supported by the National Science Foundation under Grant Nos. 1069125 and 1116360.

REFERENCES

1. Guzdial M. Human-centered Computing: a new degree for Licklider's world. Commun ACM. 2013;56(5):32-34.
2. Jaimes A, Dimitrova N. Human-centered multimedia: culture, deployment, and access. IEEE MultiMedia. 2006;13(1):12-19.
3. Panchanathan S, McDaniel T, Balasubramanian V. Person-centered accessible technologies: improved usability and adaptation through inspirations from disability research. In: Proceedings of the 2012 ACM Workshop on User Experience in e-Learning and Augmented Technologies in Education; 2012 Oct 29-Nov 2; Nara, Japan. New York (NY): ACM; 2012. p. 1-6.
4. Panchanathan S, Chakraborty S, McDaniel T, Tadayon R. Person-centered multimedia computing: a new paradigm inspired by assistive and rehabilitative applications. IEEE MultiMedia. 2016;23(3):12-19.
5. Stephanidis C. User interfaces for all: concepts, methods, and tools. Hillsdale (NJ): CRC Press; 2000.
6. Panchanthan S, Krishnan NC, Krishna S, McDaniel T, Nallure Balasubramanian V. Enriched human-centered multimedia computing through inspirations from disabilities and deficit-centered computing solutions. In: Proceedings of the 3rd ACM International Workshop on Human-centered Computing; 2008 Oct 26-31; Vancouver, Canada. New York (NY): ACM; 2008. p. 35-42.
7. Knapp ML, Hall JA, Horgan TG. Nonverbal communication in human interaction. Boston (MA): Cengage Learning; 2013.
8. Segrin C, Flora J. Poor social skills are a vulnerability factor in the development of psychosocial problems. Hum Commun Res. 2000;26(3):489-514.
9. McDaniel T, Bala S, Rosenthal J, Tadayon R, Tadayon A, Panchanathan S. Affective haptics for enhancing access to social interactions for individuals who are blind. In: Stephanisdis C, Antona M, editors. Universal Access in Human-Computer Interaction. Design and Development Methods for Universal Access. New York (NY): Springer; 2014; p. 419-429. (Lecture Notes in Computer Science).
10. McDaniel T, Krishna S, Balasubramanian V, Colbry D, Panchanathan S. Using a haptic belt to convey non-verbal communication cues during social interactions to individuals who are blind. In: 2008 IEEE International Workshop on Haptic Audio Visual Environments and Games. New York (NY): IEEE; 2008. p. 13-18.
11. Ekman P, Friesen WV. Facial action coding system: investigator's guide. Washington (DC): Consulting Psychologists Press; 1978.
12. Bala S, McDaniel T, Panchanathan S. Visual-to-tactile mapping of facial movements for enriched social interactions. In: 2014 IEEE International Symposium on Haptic, Audio and Visual Environments and Games (HAVE) Proceedings; 2014 Oct 10-11; Dallas, TX. New York (NY): IEEE 2014. p. 82-87.
13. Valstar MF, Pantic M. Biologically vs. Logic inspired encoding of facial actions and emotions in video. In: 2006 IEEE International Conference on Multimedia and Expo; 2006 Jul 9-12; Toronto, Canada. New York (NY): IEEE; 2006. p. 325-328.
14. Geldard FA, Sherrick CE. The cutaneous "rabbit": a perceptual illusion. Science. 1972;178(4057):178-179.
15. Venkateswara H, Chakraborty S, Panchanathan S. Deep-learning systems for domain adaptation in computer vision: learning transferable feature representations. IEEE Signal Process Mag. 2017;34(6):117-129.
16. Venkateswara H, Eusebio J, Chakraborty S, Panchanathan S. Deep hashing network for unsupervised domain adaptation. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2017 Jul 21-26; Honolulu, HI. New York (NY): IEEE; 2017. p. 5385-5394.
17. Venkateswara H, Lade P, Ye J, Panchanathan S. Coupled support vector machines for supervised domain adaptation. In: Proceedings of the 23rd

ACM International Conference on Multimedia; 2015 Oct 26-30; Brisbane, Australia. New York (NY): ACM; 2015. p. 1295-1298.

18. Venkateswara H, Chakraborty S, Panchanathan S. Nonlinear embedding transform for unsupervised domain adaptation. In: Computer Vision – ECCV 2016 Workshops. New York (NY): Springer; 2016. p. 451-457. (Lecture Notes in Computer Science).
19. Lucey P, Cohn JF, Kanade T, Saragih J, Ambadar Z, Matthews I. The extended Cohn-Kanade dataset (CK+): a complete dataset for action unit and emotion-specified expression. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition – Workshops; 2010 Jun 13-18; San Francisco, CA. New York (NY) IEEE: 2010. p. 94-101.
20. Pantic M, Valstar M, Rademaker R, Maat L. Web-based database for facial expression analysis. In: 2005 IEEE International Conference on Multimedia and Expo; 2005 Jul 6-8; Amsterdam, The Netherlands. New York (NY): IEEE; 2005. p. 5.
21. Godwin KM, Wasserman J, Ostwald SK. Cost associated with stroke: outpatient rehabilitative services and medication. Top Stroke Rehabil. 2011;18(sup1):676-684.
22. Bass J, Baum C, Christiansen C. Person-Environment-Occupation-Performance (PEOP): An occupation-based framework for practice. In: Hinojosa J, Kramer P, Royeen C, editors. Perspectives on human occupation: theories underlying practice. 2nd ed. Philadelphia (PA): F.A. Davis; 2017. p. 161.
23. Radomski MV, Latham CAT. Occupational therapy for physical dysfunction. Philadelphia (PA): Lippincott Williams & Wilkins; 2008.
24. Tadayon R. 2017. A person-centric design framework for at-home motor learning in serious games [dissertation]. [Tempe, AZ]: Arizona State University.
25. Tadayon R, McDaniel T, Goldberg M, Robles-Franco PM, Zia J, Laff M, Geng M, Panchanathan S. Interactive motor learning with the autonomous training assistant: a case study. In: Human-Computer Interaction: Interaction Technologies. New York (NY): Springer; 2015. p. 495-506. (Lecture Notes in Computer Science).
26. Wolf SL, Catlin PA, LinkArcher A. Assessing Wolf motor function test as outcome measure for research in patients after stroke. Stroke. 2001;32(7):1635-1639.
27. Barthel DW. Functional evaluation: the Barthel index. Med J. 1965;14:16-65.
28. Shah N, Basteris A, Amirabdollahian F. Design parameters in multimodal games for rehabilitation. Games Health J. 2013;3(1):13-20.
29. Shute VJ. Stealth assessment in computer-based games to support learning. Computer Games and Instruction. 2011;55(2):503-524.
30. Sigrist R, Rauter G, Riener R, Wolf P. Augmented visual, auditory, haptic, and multimodal feedback in motor learning: a review. Psychon Bull Rev. 2013;20(1):21-53.
31. Czisikszentmihalyi M. Flow: the psychology of optimal experience. New York (NY): Harper & Row; 1990.
32. Chen J. Flow in games (and everything else). Commun. ACM. 2007;50(4):31-34.