SIGACCESS
NEWSLETTER
A regular publication of ACM SIGACCESS: Special Interest Group on Accessible Computing

Exploring the Use of an Aerial Robot to Guide Blind Runners

Eelke Folmer

Though reducing health disparities for individuals with disabilities was identified as a research priority more than a decade ago, little progress has been made so far.

In the first article, Eelke Folmer presents an innovative project that aims to use unmanned aerial robots to guide blind runners on a track. He describes a prototype and identifies some of the research challenges his lab is trying to tackle.

An Emotion Regulation App for School Inclusion of Children with ASD: Design Principles and Preliminary Results for its Evaluation

Charles Fage

Preliminary results show that our app provides children with a relevant self-regulation support in classrooms.

Following his Best Paper Award at ASSETS'14, Charles Fage writes an article about his work with children with Autism Spectrum Disorders. We reports on a 3-month pilot study with 10 children in mainstreamed schools.

Getting Fast, Free, and Anonymous Answers to Questions Asked by People with Visual Impairments

Erin Brady

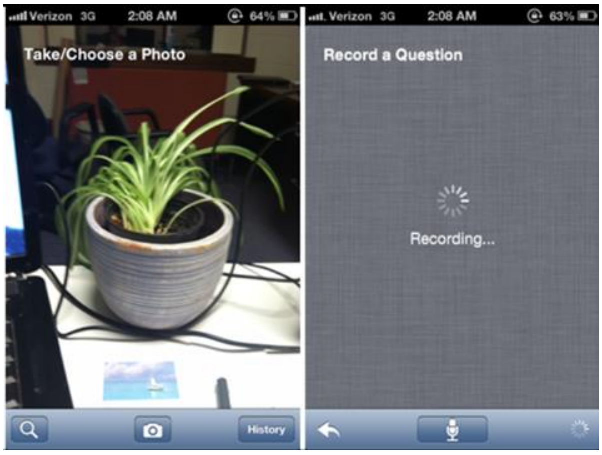

We are able to improve upon the fast speeds and quality answers achieved in VizWiz Social by combining *friendsourcing* and *crowdsourcing*, while preserving anonymity and getting answers for free.

In the third and last article, Erin Brady investigates the use of a human-powered access tool that connects people with visual impairments to sighted workers who can answer their visual questions.

Hugo Nicolau
Newsletter editor

Inside this Issue

About the Newsletter

SIGACCESS is a special interest group of ACM on Accessible Computing. The SIGACCESS Newsletter is a regular online publication of SIGACCESS that includes content of interest to the community. To join SIGACCESS, please visit our website **www.sigaccess.org**

SIGACCESS Officers

Andrew Sears
Chair

Clayton Lewis
Vice-Chair

Shari Trewin
Secretary Treasurer

Hugo Nicolau
Newsletter Editor

Joshua Hailpern
Information Director

Want to see your article featured in the next issue?

Contributions and inquiries can be emailed to **editors_sigaccess@sigaccess.acm.org**

We encourage a wide variety of contributions, such as letters to the editor, technical papers, short reports, reviews of papers or products, abstracts, book reviews, conference reports and/or announcements, interesting web page URLs, local activity reports, and so forth. Actually, we solicit almost anything of interest to our readers and community.

You may publish your work here before submitting it elsewhere. We are a very informal forum for sharing ideas with others who have common interests. Anyone interested in editing a special issue on an appropriate topic should contact the editor.

Exploring the use of an aerial robot to guide blind runners

Eelke Folmer
University of Nevada, Reno
eelke.folmer@gmail.com

Introduction

Lack of physical activity is a serious health concern for individuals with visual impairments who have fewer opportunities to engage in physical activities that provide the amounts and kinds of stimulation to maintain adequate fitness and to support a healthy standard of living. Adolescents with visual impairments often exhibit delays in motor development, such as poor balance and inefficient gait that are considered by-products of predominantly sedentary behaviors during the developmental years [1]. Physical inactivity has been associated with several medical conditions, such as cardiovascular disease, high blood pressure, cancer, type 2 diabetes, and depression.

Though reducing health disparities for individuals with disabilities was identified as a research priority more than a decade ago [2], little progress has been made so far [3]. Exercise opportunities for this population have identified to be limited due to safety concerns, lack of accessible exercise equipment and a general lack of accessible activities to choose from [4]. Ironically, efforts are underway to label obesity as a disability [5], which puts this population at risk of being even further marginalized and underserved.

A number of adapted physical activities exist, such as Goal Ball or Beep Baseball (see [6] for an overview), but these activities cannot be played independently. In recent years, accessible exercise games, such as VI Tennis [7] or Eyes-Free Yoga [8] have been developed that can be played independently and which could increase existing exercise opportunities of individuals with visual impairments. Accessible exergames have enjoyed some success, e.g., the Vifit.org website lists three accessible exergames (Tennis, Bowling, Whac-a-mole) that can be played using a PC and Wii remote have been downloaded more than 20,000 times. A number of schools of the Blind are using these games in their physical education programs. Whether exergames can engage children or adults in levels of physical activity that are high enough to be considered healthy has been a topic of debate due to contradicting studies [9] and a lack of tools that precisely measure the energy expenditure of playing an exergame. Because of this recent criticism and the fact that exercise games have decreased significantly in popularity, my lab is directing its research efforts to explore how to make existing physical activities more accessible to users who are blind.

Popular physical activities like cycling or running can be made accessible to users with visual impairments through the use of a sighted guide. Running is one of the oldest most popular physical activities that is practiced regularly by nearly 20% of the US population. Running has numerous well established health benefits that include weight loss, improved cardiovascular function, increased lifespan, and decreased effects of aging [7]. Blind individuals can run safely and independently on a treadmill, but running outside is not only healthier but also considered a more enjoyable experience. Though a guide dog can help a blind individual navigate, guide dog users are instructed not to run with their dogs, because: (1) guide dogs aren't trained to lead their users around obstacles while running; and (2) when the dog runs with a harness there is a high risk the user unintentionally provides corrective cues to the dog that modifies their behavior.

Blind athletes can run safely on a track or in an outdoor environment by having a sighted guide run to their side, which guides them using a tether rope (See figure 1). Guide runners may provide verbal cues to indicate obstacles or to indicate a change in terrain. Changes in direction, such as a bend in the road, are conveyed using the tether. A guide needs to consider the ground width of the pair of runners for collision avoidance. Urgent stops are conveyed using a short pull on the tether. The footsteps of the sighted guide may contain useful information about the terrain.

Figure 1 : A sighted guide guiding a blind athlete using a tether rope at the Paralympics

In addition to the challenge of meeting someone who has a compatible schedule and is willing to exercise together, this dependency on others puts significant constraints on the frequency and duration with which blind individuals can exercise.

Approach

It has been suggested that robots may allow individuals with disabilities to lead more independent lives [11]. To help blind individuals navigate, the use of grounded robots has been previously explored [12-14]. These systems typically involved a wheeled robot that is instrumented with various cameras and sensors, which guides the blind user using a leash (similar to a guidedog).

Because running is performed at much faster pace than walking, our project explores the use of an unmanned aerial robot, i.e., a drone, to guide a blind runner on a track. This context is interesting due to the strict time constraints to correct for veering to assure the safety of the runner. A related research project "Joggobot" [15] recently demonstrated the use of a quadrotor as an exercise companion for sighted runners. This project appropriates a commercially available quadcopter (Parrot AR quadrotor) to fly a fixed distance ahead of a runner. Where a sighted runner can easily follow a quadrotor using visual information, our research project seeks to make contributions in the area of non-visual human-robot interaction by exploring how to convey the location of the aerial robot through non-visual means. An aerial robot doesn't require an athlete to wear any sensors or hold a leash, which does not impede a blind athlete's ability to run. Compared with grounded guide robots, an aerial robot may offer the following benefits:

(1) Vantage point: Because the aerial robot flies ahead of the runner at eye level it may be able to detect obstacles earlier with a higher accuracy than a grounded robot or a wearable solution.

(2) Costs: quadrotors are much cheaper than a grounded robot, which is important, as assistive technology is often prohibitively expensive.

(3) Obstacle avoidance: Aerial robots do not have to accommodate for rough surfaces and can negotiate stairs, slopes or obstacles, rocks, potholes with no effort.

(4) Speed: Quadrotors can fly with a maximum speed of 11 mph, which fast enough to accommodate most runners.

Rather than using a leash, we will explore using the sound of the aerial robot's rotors to convey its location. This sound may be loud enough for a blind athlete to follow the aerial robot with reasonable accuracy; alternatively a speaker could be attached to the robot to produce a sound.

Prototype implementation


Figure 2: Parrot AR 2.0 Quadrotor

We developed a prototype using the Parrot AR 2.0 platform. This quadrotor is controlled using an app on an Android phone that connects to the quadrotor's wireless hotspot. An SDK is available for developing custom apps. To stay within range of the quadrotor, a blind athlete will wear the smartphone in an armband, allowing the system to be mobile. The quadrotor features a forward facing camera (1280x720) and a downward facing camera (320x240). Ultrasound sensors are used for sensing the quadrotor height from the ground. This quadrotor can fly for approximately 24 minutes. Compared with other quadrotors, a benefit of the Parrot is that it features a Styrofoam hull, which protects the runner from the spinning rotors in case of a collision. This is an important safety feature and may help blind athletes trust our approach. For this project we limit ourselves to guiding a blind runner on an oval track with no obstacles.

The location of the runner is tracked using a fiducial marker that the runner wears on their shirt. Using the known size of the visual marker, the quadrotor can estimate its distance to the runner as well as the direction the runner is facing and adjust its location and speed to maintain a fixed distance between the runner and the aerial robot. If the quadrotor loses track of the runner, a fallback strategy is employed where the quadrotor will land and idle the rotors to allow a blind athlete to approach the quadrotor, where the fiducial marker may be seen again. To protect the blind athlete from colliding with the quadrotor, it will immediately increase its altitude when it observes that the runner gets too close.

A major challenge is localizing the quadrotor. Current GPS sensors only offer a 10 feet accuracy which isn't high enough to keep the quadrotor in a running lane. Though various high accuracy GPS units are available these are also very expensive. As an alternative, we have used the quadrotor's downward facing camera to follow one of the lines on the track. A challenge here is that only a single camera feed can be fed to the mobile app, which requires us to switch between cameras every few frames.

Ongoing research & Future work

Unfortunately our campus is located within 5 miles of an airport, therefore due to FAA regulations, we have not been able to test our system on a running track on campus with blind athletes. We have been able to test our system indoors with sighted subjects, but tracking the fiducial marker is difficult due to poor lighting conditions and the quality of the camera sensors. We are waiting for the completion of the construction of a large outdoor-netted area (600x200 feet) on our campus that the FAA considers to be an indoor space. User studies will be conducted with blind athletes that will be recruited from a local chapter from the United States Associated of Blind Athletes (USABA). Some of the research questions we plan to investigate are:

1. What type of feedback is most effective (rotor sound / audio)?
2. What is the optimal distance & height for the quadrotor to fly?
3. With what accuracy can a blind athlete follow the quadrotor?

4. What are the best strategies to correct for veering?
5. Would blind athletes trust this approach?

Given a successful outcome of our project, aerial robots could be used to guide blind athletes for other athletic activities, such as swimming or cycling. Our project may allow blind athletes to run on a track safely and independently; thus increasing opportunities for this population to be physically active. It has been suggested that robots could eventually replace guide dogs. Dog guides are expensive to train and have a limited lifespan, where aerial robots are available at a relative low cost. A current limitation of aerial robots is that their flight time is limited, but this will no doubt improve over time.

Conclusion

Aerial robots are making a transition from military weapons to service based tools that deliver packages, inspect bridges and take photos. As we are entering an era where robots will be part of our daily lives, this "moonshot" project investigates how aerial robots can allow blind athletes to run independently, and develop a fundamental and transformative understanding of non-visual human-robot interaction.

Acknowledgments

This research supported by NSF Grant CHS-1445380
Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

1. James H Rimmer, Jennifer L Rowland, and Kiyoshi Yamaki. Obesity and secondary conditions in adolescents with disabilities: addressing the needs of an underserved population. J Adolesc Health, 41(3):224–9, Sep 2007.
2. Healthy people 2010: understanding and improving health. Technical report, United States. Dept.of Health and Human Services., Washington, DC, 2000.
3. Michele Capella-McDonnall. The need for health promotion for adults who are visually impaired. Journal of Visual Impairment & Blindness, 101(3):133–145, 2007.
4. Longmuir, P. E., and Bar-Or, O. Factors influencing the physical activity levels of youths with physical and sensory disabilities. Adapted Physical Activity Quarterly 17 (2000), 40–53.
5. Leahey, H. The washington post: Obesity is a disability, http://www.washingtonpost.com/posteverything/wp/2014/12/18/obesity- is- a- disability- employers- should- start- treating- it- that- way/.
6. L. Lieberman, *Adapted Physical Education and Sport*, ch. Visual Impairments. Human Kinetics, 2005.
7. Morelli, T., Foley, J., Columna, L., Lieberman, L., and Folmer, E. VI-Tennis: a vibrotactile/audio exergame for players who are visually impaired. In Proceedings of the Fifth International Conference on the Foundations of Digital Games, FDG '10, ACM (New York, NY, USA, 2010), 147–154.
8. Rector, K., Bennett, C. L., and Kientz, J. A. Eyes-free yoga: An exergame using depth cameras for blind & low vision exercise. In Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility, ASSETS '13, ACM (New York, NY, USA, 2013), 12:1–12:8.
9. Biddiss, E., and Irwin, J. Active video games to promote physical activity in children and youth: a systematic review. Arch Pediatr Adolesc Med 164, 7 (Jul 2010), 664–672.

10. E. F. Chakravarty, H. B. Hubert, V. B. Lingala, and J. F. Fries, "Reduced disability and mortality among aging runners: a 21-year longitudinal study," *Arch Intern Med*, vol. 168, pp. 1638–46, Aug 2008.

11. S. W. Brose, D. J. Weber, B. A. Salatin, G. G. Grindle, H. Wang, J. J. Vazquez, and R. A. Cooper, "The role of assistive robotics in the lives of persons with disability," American Journal of Physical Medicine & Rehabilitation, vol. 89, no. 6, 2010.

12. [21] H. Mori and S. Kotani, "Robotic travel aid for the blind: Harunobu-6," in *European Conference on Disability, Virtual Reality, and Assistive Technology*, 1998.

13. [22] V. Kulyukin, C. Gharpure, J. Nicholson, and G. Osborne, "Robot-assisted wayfinding for the visually impaired in structured indoor environments," *Autonomous Robots*, vol. 21, no. 1, pp. 29–41, 2006.

14. [23] G. Galatas, C. McMurrough, G. L. Mariottini, and F. Makedon, "eyedog: An assistive-guide robot for the visually impaired," in *Proceedings of the 4th International Conference on PErvasive Technologies*

15. F. Mueller, E. Graether, and C. Toprak, "Joggobot: jogging with a flying robot," in *CHI'13 Extended Abstracts on Human Factors in Computing Systems*, pp. 2845–2846, ACM, 2013.

About the Author:

Eelke Folmer is an associate professor at the University of Nevada, where he runs the Human+ lab (eelke.com). His lab explores the use of emerging technology, such as wearable computing and aerial robots, to augment the abilities of users with visual impairments, as to improve their daily lives, health and wellbeing.

An Emotion Regulation App for School Inclusion of Children with ASD: Design Principles and Preliminary results for its Evaluation

Charles Fage
Université de Bordeaux, Bordeaux, France
Inria Bordeaux Sud-Ouest, Talence, France
Charles.fage@inria.fr

Abstract

The inclusion of children with Autism Spectrum Disorders (ASD) in mainstreamed environments is critically impeded by their difficulties in self-regulating their emotions. We present the design and evaluation of a tablet-based application dedicated to supporting children with ASD in self-regulating their emotions. This system relies on well-proven (paper-based) emotion-regulation interventions, made by therapists and parents. To identifying both the design principles and mainstreamed-environment constraints, we adopted a participatory design. Preliminary results are presented relative to five children with ASD (and five control ones) who have used our system during three months in mainstreamed schools. They suggest that not only is our system successfully used autonomously in mainstreamed classrooms, but it is also efficient in supporting children to self-regulate their emotions, allowing them to sustain mainstream inclusion.

Introduction

Emotion self-regulation represents a challenge for children with Autism Spectrum Disorders (ASD) as they struggle to manage emotions and maintain control of behaviours [5]. Neurotypical children who demonstrate such skills are engaged in school activities and have better academic competences. Self-regulation also plays a critical role to successfully include children with ASD in a range of mainstreamed environments, such as school.

In practice, two main co-regulation approaches are used to develop emotion self-regulation, known as Emotion-Regulation Intervention (ERI). The first approach relies on well-proven, paper-based, clinical intervention methods to help individuals recognize their emotions. For an individual, recognizing his own emotional states is a fundamental cognitive ability for self-regulation. These interventions are conducted by therapists in their office, not in mainstreamed environments, limiting their impact in daily life. The second approach relies on the child's social environment such as parents develop practical and efficient coping strategies that they provide to their child with personalized support (*e.g.* favourite pictures). However, this co-regulation requires the presence of a parent.

Bringing ERI to mainstreamed environments raises difficulties since they rely on paper-based supports and/or the presence of a parent. Touchscreen tablets can be used on-the-go. Furthermore, they are increasingly used as portable gaming platforms, making them a non-stigmatizing support to bring ERI in mainstreamed environments.

To the best of our knowledge, there is no study assessing a technological support for an ERI that supports children with ASD in mainstreamed environments.

We present the design and evaluation of a tablet-based application that realizes ERI by leveraging both paper-based and parent-regulation strategies [2]. This application allows a child with ASD to autonomously recognize and manage their emotions inside the classroom. We first present classroom requirements and design principles that come from a participatory design approach, involving all stakeholders of the children. This approach was a key factor to the adoption and usage of our tool by all the stakeholders. Second, we present our tablet-based application for ERI that leverages our proposed design principles. Third, we conducted a pilot study with five students with ASD (and five control children) that used our ERI application in mainstreamed schools over a period of three months. We show that our application has been autonomously used, and reduced inappropriate behaviours, as assessed by both teachers and a school aide.

Design Principles of ERI on a Tablet-based Application

Designing a technology to be used in mainstreamed environments raises challenges. Indeed, these environments consist of a variety of people, often unaware of the specificities and needs of children with ASD. Participatory design is an approach that involves this variety of people in the design process of a technology. Participatory design involves end-users and their stakeholders to identify needs and constraints. It has been extensively used in the design of technology for children with ASD [1]. Adopting this approach, we worked with families, teachers, special-education teachers, school aides and therapists to identifying classroom usage requirements and proposing design principles to implement Emotion-Regulation Intervention on a tablet-based support.

We identified three classroom usage requirements, to ensure our ERI support will fit this environment. The first constraint is related to the classroom instructional flow, which is critical for some children, especially with ASD. School staff and therapists were unanimous on the fact that *the intervention had to be as short as possible* to prevent the child from losing track of what is going on in the classroom. Second, *ERI must not use the auditory channel*. Audio materials were excluded to conduct the intervention inside the classroom. Although headphones could be used, this would induce sensory exclusion from the classroom and stigmatization. Third, *ERI must promote reading skills*. Even such skills are not the purpose of ERI, this is a pervasive need in the school setting. Consequently, supporting this skill in any activity at school fits the school learning objectives. To support this, visual double-coding (*i.e.* pictorial and textual) has been suggested by the school staff for every interaction item of our application.

After identifying usage requirements, we propose six principles for ERI that leverage and combine concepts and strategies discussed in the literature. These principles were presented to families, school staff and therapists to ensure they matched the constraints we just discussed.

Emotion identification and naming

Many ERI focus on helping children identify their emotions and correctly naming them [7]. The procedure consists to provide the children with word-emoticon pairs of basic emotions, such as anger, fear, sadness and joy. *Principle 1: For emotion identification, word-emoticon pairs are to be displayed and the user is to be prompted to select one pair.*

Coping strategies through idiosyncratic parental support

Parents develop coping strategies, based on co-regulation, that often rely on idiosyncratic visual support. In practice, idiosyncratic supports gather personal pictures and objects specific to each child. In fact, ERI models are consistent with this approach [7] they promote families to select/build idiosyncratic supports (*i.e.* specific to a person). The materials used by parents have demonstrated their efficacy in practice; they include objects that children like, direct instructions with appropriate words, pleasant activities, *etc*. *Principle 2: Coping strategies by co-regulation are to be developed with idiosyncratic parental support, involving visual media.*

Emotional intensity level rating

When conducting an ERI, the coping strategy to be used is determined by the emotional intensity. Thermometer-like, paper-based tools are widely used to rate the intensity of the emotion [2]. *Principle 3: A thermometer-like mechanism is to be used to allow the user to select the intensity of its emotion.*

Coping strategies adapted to the emotional intensity

Parents have reported that they usually use the same coping strategies for both positive and negative emotions. The factor that drives their choice of a particular strategy is the level of emotional intensity rather than its type. *Principle 3: The emotional intensity level is to drive the selection of a coping strategy, not the type of emotion.*

A different type of media for each emotional intensity level

Parents adapt their coping strategies with respect to the intensity level of the emotion. This approach is consistent with most ERIs [7]. A low-emotional intensity level can be addressed with relaxing instructions. However, motor activity, such as walking around the house before going to school, is required to cope with a high-emotional intensity level. *Principle 5: The effect of the media contents on the child has to match the intensity level of their emotions (e.g. non-idiosyncratic photos, idiosyncratic photos and videos).*

Step-based regulation strategy

Step-based strategies are widely used with children with ASD to perform activities. They render the environment predictable and safe, and they are known to reduce anxiety for this population [2,3]. *Principle 6: Regulation strategies are to be delivered in a step-based navigation fashion.*

Application Description

In this section, our design principles are implemented in a tablet-based application, fitting the classroom usage requirements. We selected a touchscreen tablet to run our application. This platform enables rich visual supports and allows the application to be used in any environment. Furthermore, tablets do not carry any stigma. Their effectiveness to support intervention has already been demonstrated in the context of children with ASD [4].

Interface Description: Two-Step Intervention

As suggested by ERI models [7], our intervention is structured into two steps (Principle 6): 1) emotion identification and 2) co-regulation strategies.

Step 1: Emotion identification

First, the student is invited to identify their emotion by clicking on emoticons (Principle 1). The number of available emotions has been voluntarily reduced to meet the teachers' requirements. Note that positive and negative emotions have not been separated, as children with ASD can be overwhelmed by any emotion (see Figure 1).

Then the student selects a level of emotional intensity (Principle 3) via a thermometer with a scale from 1 to 4, as advised by Gagné *et al.* [2]. The selection of an intensity level is possible only after selecting an emoticon, to reinforce the structure of the interaction (Principle 6). When a level of emotional intensity is selected, the application displays a new screen dedicated to delivering co-regulation strategies.

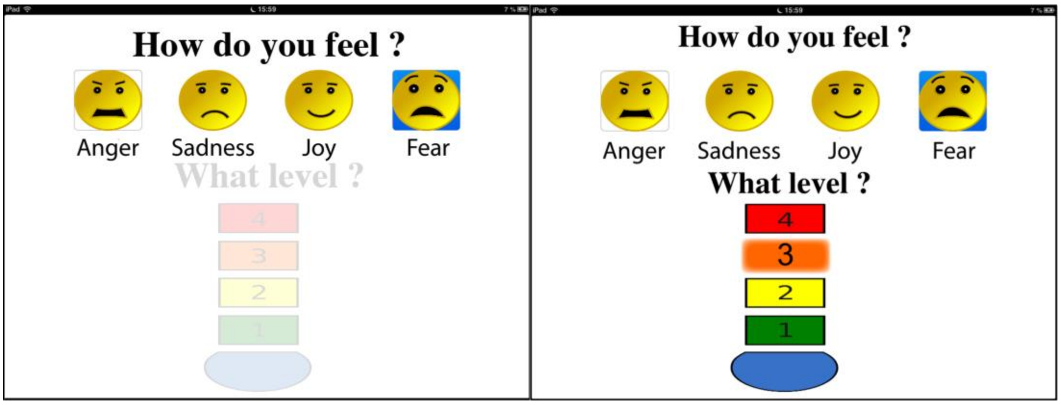

Figure 1 First step: emotion identification

Step 2: Co-regulation strategies

Associated with each intensity level is a different coping strategy (Principle 4) that takes the form of either respiratory relaxation or idiosyncratic multimedia contents (Principle 2). ERI models promote families to select/build idiosyncratic supports (*i.e.* specific to a person) [7]. In practice, idiosyncratic supports gather personal pictures and objects specific to each child. These idiosyncratic visual supports usually present the child, either in an activity they like or standing in a reassuring place. Each intensity level corresponds to a specific coping strategy (Principle 5).

Level "1" offers relaxation methods (*e.g.* respiratory) through a three-step slideshow. A textual statement and a picture illustrate every step. Relaxation methods are widely used by therapists and special education teachers to support children with ASD [2].

Level "2" provides a co-regulation strategy that consists of displaying an idiosyncratic library of photos. The child can control the pace of the slideshow with "forward/backward" buttons.

Level "3" also delivers a parent-based, co-regulation strategy, but in the form of a personalized video. A switch button is also displayed with a "Sound" label - default is set to "0" (*i.e.* mute) - to fit the classroom usage requirements from our participatory design.

Level "4" proposes the child to leave the classroom and relax in a quiet place (usually the special classroom or the library) for a duration to be determined with them. According to school staff, it is a common practice when children simply cannot self-regulate their emotions.

As can be noticed, by design, our ERI application requires to be customized by idiosyncratic visual materials selected by families and endorsed by their child. This design allows flexibility and adaptability of the materials.

Evaluation

Our pilot study took place in special education classrooms located in secondary schools. To assess the benefits of our tablet-based application, a key constraint had to be taken into account. Because there is no equivalent paper-based system, we could not compare our system with a paper-based support. It would have been cumbersome for the children to manipulate and navigate such support; as well, its nature would have excluded visual support equivalent to a tablet (*i.e.* personalized videos).

Participants

Ten students between the ages of 13 and 16 were included in our study. Five of them were children with ASD and the five others were children without ASD. Children with ASD in our study are on the lower end of the spectrum of intellectual functioning. The two groups were matched with the chronological age and the intellectual functioning (according to the IQs estimated from abbreviated WISC-IV). Neuropediatricians examined all the children and the ASD diagnosis was performed according to the criteria of the DSM-IV and with respect to the "Autism Diagnostic Interview-Revised" scale.

Materials and Instruments

We measured: 1) the application usability, and 2) the application efficacy as an emotional coping device in mainstreamed classrooms.

1) Application usability.

Child usage. At the end of each month of intervention, the school aide was asked to indicate whether the child used the application in full autonomy and in adequate manner (scored 1), or whether they had needed help to use it (scored 0) in appropriate situations (*i.e.* emotional outbursts).

Parents' point of view. We selected in the USE questionnaire sub-scales assessing the parent's perceptions in terms of usability and ease of learning (with a Likert scale from 0 to 4 for each of the 15 questionnaire items), with a maximum score of 60.

2) Application efficacy in the inclusion classroom.

To measure changes in self-regulation behaviours in low-functioning individuals, the French school version of the Quebec adaptive behaviour scale (EQCA-VS) was completed by the school aide for each child after each month of our application usage. Part 2 of this questionnaire was selected because it assesses inappropriate behaviours such as violent, withdrawal and antisocial behaviours in school environments. Each sub-scale item is scored from 0 (not observed) to 3 (severely observed) with a maximum global score of 40 points.

Procedure

Prior to our intervention, we held a meeting with the inclusion teachers, the special education teacher, the school aide, the parents, and the children. The goal was to give them an overview of our procedures, to explain the importance of using our application on a regular basis in a synergistic manner, and to answer all their questions. We also gave a demonstration of our tool, explaining its functioning. Parents were asked to choose "around ten photos or pictures and a short video that soothes for [their] child" to create/identify idiosyncratic media contents to personalize the application.

Baseline: pre-intervention tests

At the baseline assessment session, the children completed the abbreviated WISC-IV. School aide completed self-regulation questionnaire (EQCA-VS) based on its observations inside the classroom.

Intervention: application usage

The participants were observed during their inclusion in the classroom (French, mathematics, history, geography, or biology). In the context of our intervention, each participant attended a new class where new situations could occur. It was a one-hour class that occurred once a week during a period of three months. A school aide, dedicated to our study, accompanied each child during inclusion. The school aide was trained to support students with disabilities, especially children with ASD. In addition, she was instructed how to use the application to play the role of a social support during inclusion.

Post-intervention tests

All other post-test measures were completed within a week after the end of the three-month intervention. All interviews were conducted at school or at home. They addressed application usability (USE questionnaire) and application effectiveness (EQCA-VS questionnaire). We also discussed how all stakeholders experienced our project, and what could be improved.

Results

Application Usability: Overall, the usability measures revealed that our emotion-regulation application is perceived as easy to use by parents and their child. Its self-initiated use is acquired within the three months of our intervention.

Application efficacy in the inclusion classroom: The main results showed that non-adaptive behaviours were significantly more prevalent in children with ASD compared to the other children at the beginning of the intervention. By contrast, by the end of the intervention, these non-adaptive behaviours disappeared in children with ASD, making the group-related differences no longer significant (see Figure 2) [Interaction effect between group * time, $p<.003$]).

The pairwise comparisons revealed that the group difference continually decreased over the course of the intervention: the first month, the second month, and the third month.

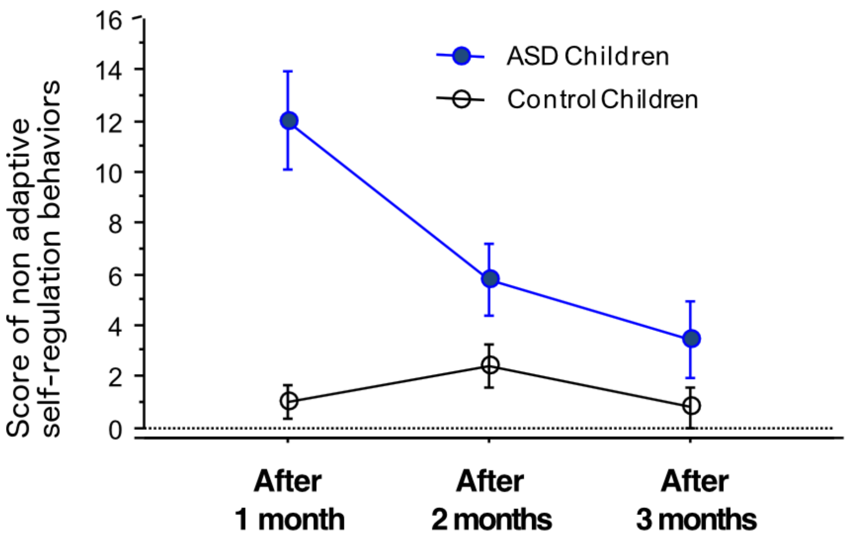

Figure 2 Number of Non-adaptive self-regulation behaviors in ASD and control children across the three-month intervention.

Discussion

To the best of our knowledge, there is no study assessing a technological support for an ERI dedicated to children with ASD in mainstreamed environments. Additionally, we found no study addressing the relevance of idiosyncratic parental co-regulations in the context of school settings to support children with ASD. The results presented here provide insights on these issues.

Our preliminary empirical results support that our ERI application provides children, especially those with ASD, with a relevant self-regulation support in mainstreamed environments, such as a classroom. Importantly, we observe that the non-adaptive behaviours in class were greatly reduced for children with ASD, despite the short intervention time (*i.e.* 3 months).

Parents reported high usability of our application. School aide reported the first step of ERI (*i.e.* emotion identification) was well structured. Within one screen and only two user-pointing inputs, our application allows children to quickly access soothing material. As a result, interaction duration maximizes the child's presence in the classroom. The simplicity and the structured interaction with the application may be factors that contributed to the fact that the children rapidly used the application autonomously.

Introducing a technological assistive tool in a mainstream classroom raises challenges. Teachers were concerned that children with ASD would spend their time playing with the tablet and others student would be distracted. School aides were concerned about not being able to handle the tool. Families were concerned about fatigue incurred by their child in a stressful environment. However, by involving all these stakeholders early in the design process, we overcame those concerns and ensured the infusion of our technological support in the school environment.

The collaborative nature of our intervention allowed our tool to be pervasively used by all stakeholders of the child's mainstreamed environment: the school. As well, we argue that participatory design allowed our tool to be appropriately used by the school staff, who had never used this technology before. Every stakeholder was able to help children using it during the first month of the intervention, and to supervise them for the last month. Leveraging the observed collaborative usage of our application, we could adapt the interface to bring the child to find someone around to help co-regulating.

Limitations and Future Work.

Regarding the participating children, their number did not reach a sample size that is required to actually have statistical conclusive results, even though the use of non-parametric statistical tests has been respected. We currently perform a study with a sample's size of children with ASD statistically reliable ($n>40$) divided into two groups (equipped vs. not equipped). Additionally, children with intellectual disabilities are also included. Such study design will enable to capture ERI app impact among children with ASD (Comparisons between equipped and not equipped children with ASD), but also its ASD sensibility (comparisons between equipped children with ASD and equipped children with intellectual disabilities (*i.e.* without ASD)).

Importantly, besides helping to regulate emotions, our application collected data regarding its usage (*i.e.* number of uses in the inclusion class, types and levels of emotion). Processing such data will help us to know whether children with ASD mostly used our application when they experienced emotional states that were critically intense or not. We would then be able to link the intensity to the type of media contents used, assessing relevance of idiosyncratic contents.

Conclusion

This paper presents a tablet application to support emotion regulation of children with ASD in mainstreamed environments. With a participatory design, we identified usage requirements, and design principles that allowed our application to be infused in a mainstreamed environment: the school. Using these design principles, other applications could be implemented to offer more adaptability to closely match the needs of children with ASD. This application has been used by ten children with and without ASD during their inclusion in secondary schools. All children successfully adopted our application thanks to its structured interface and its idiosyncratic contents. Using our application allowed children with ASD to better self-regulate their emotions.

Acknowledgments

We are deeply thankful to school staff members, families and children we worked with. We also want to thank neuropediatricians from Centre de Ressource Autisme of Bordeaux for their precious help and French Education Ministery for opening access to classrooms and supporting this project.

References

16. Frauenberger, C., Good, J., Alcorn, A., and Pain, H. Supporting the design contributions of children with autism spectrum conditions. In *Proceedings of the 11th International Conference on Interaction Design and Children*, IDC '12, ACM (New York, NY, USA, 2012), 134–143.

17. Gagné, A. *Including a Student with Autism Spectrum Disorder in my Classroom*. Conseil des coles publiques de l'Est de l'Ontario, (In French), 2010.

18. Hayes, G. R., Hirano, S., Marcu, G., Monibi, M., Nguyen, D. H., and Yeganyan, M. Interactive visual supports for children with autism. *Personal and ubiquitous computing* 14, 7 (2010), 663–680.

19. Hourcade, J. P., Williams, S. R., Miller, E. A., Huebner, K. E., and Liang, L. J. Evaluation of tablet apps to encourage social interaction in children with autism spectrum disorders. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, ACM (New York, NY, USA, 2013), 3197–3206.

20. Konstantareas, M. M., and Stewart, K. Affect regulation and temperament in children with autism spectrum disorder. *Journal of Autism and Developmental Disorders 36*, 2 (2006), 143–154.

21. Park, J. H., Abirached, B., and Zhang, Y. A framework for designing assistive technologies for teaching children with asds emotions. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '12, ACM (New York, NY, USA, 2012), 2423–2428.

22. Prizant, B. M., Wetherby, A. M., Rubin, E., and Laurent, A. C. The scerts model: A transactional, family-centered approach to enhancing communication and socioemotional abilities of children with autism spectrum disorder. *Infants & Young Children 16*, 4 (2003), 296–316.

About the Author:

Charles Fage is a third year PhD student in cognitive sciences at the University of Bordeaux in collaboration with Inria Bordeaux Sud-Ouest center. His research interest includes human centred computing, assistive technologies, autism spectrum disorders, and school inclusion. For his thesis research, he focuses on assistive applications to promote school inclusion of children with ASD and children with intellectual disabilities in mainstream classrooms.

GETTING FAST, FREE, AND ANONYMOUS ANSWERS TO QUESTIONS ASKED BY PEOPLE WITH VISUAL IMPAIRMENTS

Erin Brady
University of Rochester
brady@cs.rochester.edu

Introduction

Much of the information we encounter in daily life is visual. For people with visual impairments, this information is not inherently accessible. Technology can be used to provide access to this visual information, through automated tools like currency identifiers, object recognizers, and optical character recognition, which allow for immediate access to visual information in a users' environment. However, these tools are often are limited to a specific domain, and can fail in non-ideal conditions (*e.g.*, an optical character recognition tool trying to identify handwritten text).

By incorporating people into the information access pipeline, human-powered tools can allow for more complex information to be extracted from the user's surroundings. Human-powered access tools connect people who have disabilities to other people who can access information on their behalf. While people may not be as quick as automated tools, the person answering can make inferences based on prior knowledge and experiences, ask for clarification, and reason over the information provided.

In this article, we explore the long-term public deployment and lessons learned from VizWiz Social, a human-powered access tool that connects people with visual impairments to sighted workers or friends and family members who can answer their visual questions. By analyzing the use of this service, and running controlled experiments to determine how users value different answer sources, we have been able to build upon the original VizWiz design to create *social microvolunteering*, a method for getting fast, free, and anonymous answers to visual questions.

Related Work

VizWiz Social drew inspiration from two distinct areas of related work: the design of tools for people with visual impairments or other disabilities, and workflows that incorporate crowds of humans through crowdsourcing or friendsourcing.

Tools for People with Visual Impairments

There are many automated tools which provide alternative forms of access to information that is typically presented visually. For digital information, screenreaders can take the visual representation of a webpage or program and translate it into an aural or tactile format [1]. Automated mobile tools can provide access to visual information in the physical world, ranging from standalone devices like the KNFBReader (a specialty device that provided access to printed text) to automated smartphone applications (like currency identifiers, color recognizers, barcode scanners, and optical character recognition).

Not all visual tasks can be completely automated, either due to technical limitations or the complexity of the tasks. For these problems, people with visual impairments typically ask a sighted person around them for assistance. Mobile tools now allow these users to recruit sighted help remotely, increasing independence and safety [9].

Tools which utilize humans to provide access can be considered human powered access tools [4]. By inserting humans into systems where technological solutions are not feasible, these tools can provide more access to visual information than automated technologies or people alone. One of the most well-known human-powered access tools was Social Accessibility, a system where users with visual impairments could request image descriptions from people remotely [17].

Crowdsourcing and Friendsourcing

We can look to two interconnected areas of HCI research to find people to perform tasks for human-powered access tools: *crowdsourcing* and *friendsourcing*.

Crowdsourcing takes small tasks and distributes them to remote workers through online microtask platforms [16]. These workers can perform one-off tasks, or be integrated into complex workflows, where human work is inserted into algorithmic structures [12]. A large amount of existing crowd research has been performed on Amazon Mechanical Turk, an open microtask platform where workers can search for tasks (like recognizing objects, transcribing audio, or describing images) or surveys to complete [10].

While crowdsourcing is one way of recruiting humans who are able to answer questions or perform microtasks, another is *friendsourcing* these duties to members of your online social network. This can be done informally – by asking friends on Facebook for answers to questions [14], favors [8], or opinions about a topic [15] – or in more structured settings, like the Collabio social tagging project [2]. Prior research has found that answers from friends are perceived to be more trustworthy [13], though answers may be slower than desired [18].

VizWiz

To examine how crowdsourcing could be useful for visual question answering, Bigham *et al.* created a pilot application called VizWiz [3]. VizWiz was a Wizard-of-Oz mobile smartphone application for users with visual impairments. The pilot users could ask visual questions by taking photographs of their environment and recording a question about the picture. These questions would then be posted to Mechanical Turk and answered by an anonymous worker.

Using the traditional Mechanical Turk interface, the length of time to get an answer to each question might be prohibitative for a user. Using this method, answers would be received after the photograph was taken, the question was recorded, a task was posted to Mechanical Turk, a worker found and accepted the task, and the worker listened to the question and answered it. However, [3] introduced a novel technique called *quikTurkit*, which allows Mechanical Turk requesters to pre-recruit workers in order to get answers faster and with more consistent response times. For VizWiz, *quikTurkit* posts a task when the user opens the camera in the VizWiz application, so workers can be recruited and retained in parallel as the user takes their photograph and records their question. When the question is ready, workers are already available to answer, meaning the questions can be answered in **nearly realtime**.

In a pilot study with 11 users, where users were encouraged to ask at least one question a day throughout the week-long pilot, 87% of the 82 questions asked received a *good-faith* first answer – either the correct answer if the question could be answered from the photograph, or feedback explaining that the question couldn't be answered or how to fix it if the question was not

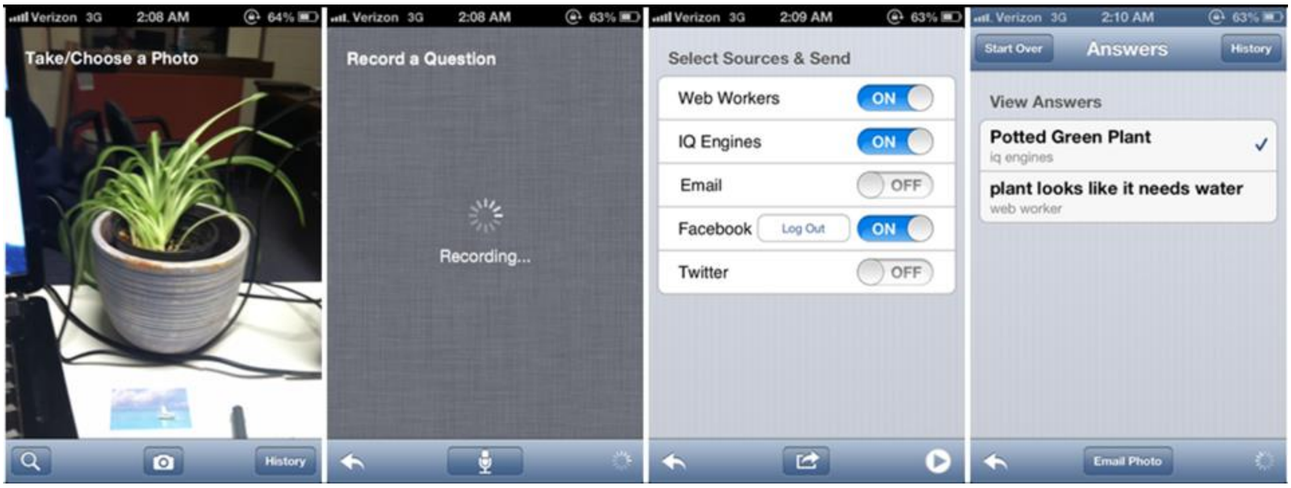

Figure 1. A screenshot of the VizWiz Social interface, where a user (a) takes a photograph, (b) records an audio question, (c) chooses the sources to send their question to, and (d) can review answers as they arrive.

answerable. By redundantly recruiting three answers to each question, each question eventually received a good-faith answer. Additionally, by pre-recruiting workers, the average time until a first answer for any question was 133 seconds. User feedback was also overwhelmingly positive, with users enjoying the service and wanting to continue to use it beyond the study period.

VizWiz Social

Based on the pilot study by Bigham *et al*., we built and deployed VizWiz Social, a live deployment of VizWiz which was released free of charge for iPhone users. Similar to the VizWiz pilot study, the user takes a photograph of visual information around them, and records a question about it. In Vizwiz Social, the user can send these questions to anonymous crowd workers, but also has the option to send the question to sighted friends and family members via email, Facebook, or Twitter, or to an automated object recognition service (Figure 1).

Since VizWiz Social was released in May 2011 to the iOS App Store, it has been used by more than 10,000 people to answer over 70,000 visual questions. By analyzing the types of questions asked with the system [5] and exploring user motivations for routing questions to specific answer sources [7], we have been able to develop social microvolunteering, a new crowdsourcing method to answer questions in a way that is fast, free, and anonymous [6].

Questions Asked with VizWiz Social

The dataset of over 70,000 questions asked with VizWiz Social since its release provides a valuable resource for researchers who want to build tools for people with visual impairments. By analyzing the questions asked, as well as features of the photographs and where those questions were sent, we can learn how to design future human-powered access tools and how to improve automated tools (like photo-taking applications).

We analyzed a sample of 1000 VizWiz Social questions to understand the way the system was being used. These questions were a randomly selected sample from the VizWiz Social questions asked in the first year of its release by users who agreed to have their data used for research purposes.

We created a taxonomy of VizWiz Social question types through an affinity diagramming process, where questions were grouped together intuitively by a team of researchers and then refined into four categories of questions, each with subcategories. The four main categories of questions were:

• **Identification questions (41%)**, where the user asked for the name or type of something pictured;

• **Reading questions (17%)**, where the user asked for a transcription of a full block of text, or for a specific piece of information from a greater section of text contained in a photograph;

• **Description questions (24%)**, where the user asked for physical qualities of something in the photograph, such as color or size;

• **Other questions (17%)**, which could not be answered due to problems with the audio or photography quality, were outside the scope of VizWiz Social, or were about the service itself.

We also looked at features of the photographs and audio recordings from users. We examined the objectivity of questions, errors in photographs (like blur or lighting issues), perceived urgency of questions, and the primary subject of the photograph for each of the 1000 questions in the sample. While there was a wide variety in the types of questions asked, the questions were typically objective (78%) and seemingly urgent (68%) questions about objects (76%), with one or two photographic errors exhibited (62%).

In addition to examining question types, we looked at user behaviors of 100 first-time users of VizWiz Social, and the 25 most active users. First time users primarily asked identification question (44%), and half of first-time users converted into multi-week users. Users with errors in their photograph were less likely to return to the service after the first day than those without errors. The 25 active users differed from first-time users, and showed changes in their usage behaviors, asking more reading questions (46%) in their later use than identification questions (25%), and showing improvement in photograph quality.

For more information about the analysis of question categories and subcategories, question features, or user behaviors, see the full paper [5].

What did we learn?

By analyzing the questions asked with VizWiz Social, we are able to draw conclusions about how to design human-powered access tools for visual question answering. We saw that most questions were objective identification or reading questions about objects, rather than more subjective description questions. As a result, answers likely do not need to be tailored to the user, and do not require much expertise, and a variety of answer sources could be combined to get high-quality answers quickly.

Answer speed and quality are important, both for guaranteeing that the tool is useful for real-life situations, and for ensuring user retention. Many questions were judged by the research team to appear *urgent* and need an answer within ten minutes. Users were less likely to return to VizWiz if the first answer to the first question they asked was not in good-faith. Future human-powered access tools should prioritize response speeds without compromising quality of responses.

Usage patterns changed over time, with the power users asking more complex questions in their later use of the service than their earlier use. Human-powered tools may be more robust to this type of evolving use than automated tools, as human workers can adapt to any reasonable question that is sent in, while automated tools often are trained for specific problem domains.

Comparing Crowdsourcing and Friendsourcing for Visual Questions

While examining the VizWiz Social dataset, we also investigated where questions were sent. Since VizWiz Social included friendsourced and automatic question answering as supplements to the crowdsourced answers in the original VizWiz pilot, we expected to see a variety of questions sent to friendsourced or automatic sources by users of the system, and wanted to learn which types of questions were more appropriate for anonymous answer sources (crowdsourced or automatic answers), and which were more appropriate for friendsourcing. Prior work had shown that friendsourcing can provide trustworthy and high-quality answers [13], and the free nature of asking questions to friends might be beneficial for users who do not wish to pay for answers, or who may need to take several photographs to fully frame the objects their question is about.

However, examination of the VizWiz Social questions revealed that very few questions were being friendsourced. Only 5% of questions in a selected month were sent to email, Facebook, or Twitter, and only 15% of users had ever tried friendsourcing their VizWiz questions. When users did try to friendsourcing questions, they had low response rates, with only 3% of questions receiving answers at all, and response times for answered questions averaging near three hours.

We sought to investigate why friendsourcing was underused for VizWiz Social questions. Did VizWiz Social users not have active Facebook or Twitter networks, or not like to ask questions of their online social networks? Were response rates and speeds so low that users did not think friendsourced answers were useful? Or were there other reasons that users preferred crowdsourced answers over friendsourced answers, such as anonymity or the fact that the users did not bear the true costs of crowdsourcing themselves?

To explore the relative values of crowdsourcing and friendsourcing for VizWiz Social users, we performed a two-part study [7]. In the first part, we surveyed the use of social network sites by people with visual impairments, to learn how they used these platforms and if they were appropriate venues for questions. In the second part, we performed an experimental study with active VizWiz Social users, testing a financial incentive for asking friendsourced questions instead of crowdsourced questions.

Survey of Social Network Site Use by People with Visual Impairments

We first surveyed the use of social network sites by people with visual impairments. We sent a link to an accessible online survey to several organizational mailing lists for people with visual impairments (including the National Federation of the Blind, American Foundation for the Blind, and the Seattle Lighthouse for the Blind) in winter 2012. This survey asked respondents about their use of social networking sites, their posting activity on the sites they used, and their thoughts about the appropriateness of platforms like Facebook and Twitter for asking questions. 191 of 203 people who responded to the online survey self-indicated that they had some level of visual impairment, and we examined results only from these 191 respondents.

The survey revealed a high level of presence on social network sites. 92% of respondents had accounts on one or more social network sites, with Facebook (80%) and Twitter (52%) being the most popularly used. While respondents often logged into the platforms, they infrequently posted statuses or new content on the sites, and even more rarely asked questions.

Respondents also indicated issues around asking questions on social network sites. 55% of respondents thought that asking questions to their online social network could be an effective way to get answers, and only 37% of Facebook users and 54% of Twitter users felt very or somewhat comfortable posting questions to the platforms they used. Users cited **technical reasons** for not using the sites for question asking – screen reader accessibility issues, 140-character limits on tweets – but also **personal reasons**, like fear of incurring social costs by annoying their networks with questions, or a preference for anonymity when asking questions. Though the survey did not ask respondents about mobile question asking tools like VizWiz, one respondent brought the application up organically, saying:

> *As for something like Vizwiz I would not always feel comfortable having my Twitter followers see what I am scanning, or say if I did not take the picture correctly. I would prefer anonymous web workers or a pre-selected group of friends.*

Exploring the Value of Friendsourcing Answers to VizWiz Social Questions

The survey of social network site use revealed that users had smaller-than-average network sizes, meaning fewer friends or followers are available at any point to answer their questions. While this alone might have been responsible for the low use of social sources in VizWiz Social, we wanted to delve deeper into users' motivations for routing questions by utilizing an artificial financial incentive to determine the value of friendsourced answers when compared to crowdsourcing.

We performed a user study with 23 active VizWiz Social users in spring 2012. Participation in the study was completely optional for users, who were presented with the details of the study when they opened the application, and could decline and go immediate to regular use of the application, or join and receive further details of the month-long study.

Each participant was compensated with $10 for participating, and given a starting bonus of $25. Each question they sent to crowdsourced or anonymous answer sources cost a nominal amount chosen to represent the true cost of crowdsourced or automated questions (which are currently absorbed by our research team). This amount was subtracted from the bonus, rather than being paid directly by the participant, so participants could earn between $10 and $35 based on how many questions they crowdsourced.

While there was no statistically significant difference in the number of questions the users asked during the month before the study and the month of the study, there were surprising uses of the different answer sources. Crowdsourcing of questions remained relatively static (81% of questions in the month before the study were crowdsourced, 83% during), and social sources remained unutilized (14% of questions in the month before the study were friendsourced, 1% during). The only significant drop was in questions sent to automatic object recognition, which dropped from 93% in the month before the study to 45% during the study.

In a survey after the study, all respondents indicated a preference for sending their questions to crowdsourcing answers sources, rather than friendsourcing, or had no preference between the two. Again, their responses had two themes: **technical reasons** for preferring crowdsourcing, because of the speed and constant availability of the workers, and **personal reasons** of wanting anonymity in their question asking.

For more information about this study or the survey of social network use by people with visual impairments, see the full paper [7].

What did we learn?

Users had both practical and personal reasons for preferring crowdsourcing to friendsourcing. While many of the technical reasons could be overcome by using different platforms or more expansive targeting of friends on social networks, the personal reasons all related to a desire for anonymity or more nuanced sharing of questions, indicating that traditional friendsourcing is unlikely to be a good resource for human-powered access tools.

When compared to general online adults, our respondents' presence on social network sites accounts appears higher than average, though responses may be biased due to recruitment through online mailing lists. However, on average users had fewer friends or followers on the platforms, which may contribute to their lower response rates and speeds when asking questions on Facebook or Twitter. While we did not investigate this further for our study, later research on the Facebook platform revealed that screenreader users had smaller network sizes than average Facebook users, as the site became accessible after it was released, and users with visual impairments started their accounts later than comparable non-screenreader users [19]. These differences may even out over time, but accessibility is often added into sites in later phases of development [11], and similar issues are likely to arise with future social network sites.

Social Microvolunteering

Our investigation of the questions asked with VizWiz Social revealed that most were objective and appeared urgent, meaning that the speed of answers was more important than the source of workers. By exploring use of social network sites for question asking, we found that both VizWiz Social users and general social network users with visual impairments did not want to use their social networks for question asking, preferring crowdsourced answers which had better response rates, speeds, and were anonymous.

However, the financial costs of crowdsourcing cannot be absorbed by the research team indefinitely, as our work on VizWiz Social has been funded by research grants. While this is sustainable in the short term, it cannot generalize to other human-powered access tools, and the costs may prove prohibitive if passed directly on to users.

In order to address these concerns while addressing VizWiz users' desire for anonymity, we developed **social microvolunteering**, a hybrid method combining benefits of both crowdsourcing and friendsourcing to answer altruistic tasks. In social microvolunteering, a core group of volunteers have tasks for other people posted to their social network accounts. If the volunteer is online, they can complete the work; if they are not, their friends and family members can see the work and complete if it they are interested, expanding the pool of potential workers to allow for faster results.

To explore social microvolunteering in the context of visual questions, we created Visual Answers, a Facebook application to answer VizWiz-style questions. Volunteers can install Visual Answers to their Facebook account. When a VizWiz user sends in a question, it is posted to the volunteer's Facebook Newsfeed, where either the volunteer or any of their friends could answer the question in the comments. Those answers are then forwarded to the VizWiz user.

We performed a survey and a pilot study with potential Visual Answers volunteers to test the feasibility of the idea. In the survey, with 350 Facebook users, 55% of respondents said they would want to use an application like Visual Answers, which was presented as a hypothetical application to represent the concept of social microvolunteering. The respondents who wanted to use such a social microvolunteering application thought it would be an effective way to help people with visual impairments and raise awareness of disability issues, while respondents who didn't want to use the application cited privacy concerns and uncertainty about the feasibility of the application.

After the survey, we offered all 350 participants the chance to pilot the Visual Answers application. 91 volunteers joined a 12-day pilot study, allowing Visual Answers to post 1,130 questions to their Facebook Newsfeeds. We examined both the technical metrics of the system and the volunteer reactions to Visual Answers.

While only 42% of questions received a comment, and the average first response time for any individual post was 58 minutes, response rates and speeds could be dramatically improved by posting questions to multiple volunteers at once. When questions were posted to 10 users at once, all received at least one comment and the average first response time was 4 minutes and 48 seconds. Adding additional volunteers kept decreasing the time to first response, and if a question was posted to every available volunteer, an answer could be received within seconds due to the large network of volunteers and friends who might be online. Answers were also high-quality, with 82% of all comments being good-faith answers, and 91% of the first comments on each post being good-faith answers.

Post-study surveys revealed that volunteers liked using the application. 95% of respondents felt positive about using Visual Answers, 90% thought Facebook was a good venue for social microvolunteering, and 83% wanted to use Facebook for social microvolunteering in the future.

For more information about social microvolunteering or Visual Answers, see the full paper [6]

What did we learn?

By combining friendsourcing and crowdsourcing into social microvolunteering, we are able to retain and improve upon the fast speeds and high-quality answers achieved in VizWiz Social, preserve the anonymity between users and answerers, and get answers for free from volunteers, who may gain enjoyment or feel altruistic through their participation.

In the initial study, we asked about respondents' engagement in volunteering. 51% said that they wanted to volunteer more, but lacked the time, resources, or physical ability to participate in in-person volunteer opportunities. Applications like Visual Answers may fill a need for people who want to participate further, but cannot volunteer on a fixed schedule or give much time or money.

Of the 91 volunteers who installed Visual Answers, 64% had indicated in the original survey that they had friends or family members with disabilities. For these participants, using a social microvolunteering tool may provide benefits for both the volunteer (who can make an impact for a cause they care about) and the VizWiz Social user who benefits from their work (as the volunteer may have more domain knowledge and provide higher-quality answers).

Discussion

In the four years since VizWiz Social was released, the domain of visual question answering has seen a number of advancements. Automated tools have improved, meaning that more questions can be answered on-demand. Other human-powered visual question answer tools have been released as well – TapTapSee[1] combines automated and human object recognition, and

[1] taptapseeapp.com

BeMyEyes[2] connects users with visual impairments to sighted volunteers via video stream to allow them to answer questions in real time. Despite these advancements, VizWiz Social continues to be used by the public to get answers to visual questions, and the lessons learned from VizWiz Social and Visual Answers can be extended to new human-powered access tools.

We can use the database of visual questions asked by VizWiz Social users to improve multiple aspects of the access technology field. Automated tools can use the questions and metadata from VizWiz Social in order to train better algorithms, which could facilitate better non-visual photography or perform automatic routing of questions. Designers can learn lessons from the users' selection of sources and longitudinal use of the service, in order to design access technologies which match users' expectations.

While visual questions like the ones sent to VizWiz Social are ideal for social microvolunteering, the concept could be extended and generalized to other microtasks for people with disabilities, or other altruistic domains like citizen science. There remains room for further exploration of the benefits of combining the crowdsourcing and friendsourcing, both in human-powered access tools and beyond.

Acknowledgments

The projects presented above were done in collaboration with Meredith Ringel Morris of Microsoft Research and Jeffrey P. Bigham of Carnegie Mellon University. This work has been supported by a Provost Fellowship from the University of Rochester and National Science Foundation Awards #IIS-1149709 and #IIS-1049080.

References

23. 1. Asakawa, C. and Itoh, T. User Interface of a Home Page Reader. *Proceedings of the ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*, (1998).

24. 2. Bernstein, M., Tan, D., Smith, G., Czerwinski, M., and Horvitz, E. Collabio: A Game for Annotating People within Social Networks. *UIST 2009*, (2009).

25. 3. Bigham, J.P., Jayant, C., Ji, H., et al. VizWiz: Nearly Real-time Answers to Visual Questions. *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST)*, (2010).

26. 4. Bigham, J.P., Ladner, R.E., and Borodin, Y. The Design of Human-Powered Access Technology. *ASSETS 2011*, (2011).

27. 5. Brady, E., Morris, M.R., Zhong, Y., White, S., and Bigham, J.P. Visual Challenges in the Everyday Lives of Blind People. *Proceedings of the 31st ACM SIGCHI Conference on Human Factors in Computing System (CHI)*, (2013).

28. 6. Brady, E., Morris, M.R., and Bigham, J.P. Gauging Receptiveness to Social Microvolunteering. *Proceedings of the 33rd ACM SIGCHI Conference on Human Factors in Computing System (CHI)*, (2015), (to appear).

29. 7. Brady, E., Zhong, Y., Morris, M.R., and Bigham, J.P. Investigating the appropriateness of social network question asking as a resource for blind users. *Proceedings of the 2013 ACM Conference on Computer Supported Cooperative Work (CSCW)*, (2013), 1225.

---

[2] Bemyeyes.org

30. 8. Jung, Y., Gray, R., Lampe, C., and Ellison, N.B. Favors from Facebook Friends: Unpacking Dimensions of Social Capital. *Proceedings of the 2013 ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, (2013).

31. 9. Kane, S.K., Jayant, C., Wobbrock, J.O., and Ladner, R.E. Freedom to Roam: A Study of Mobile Device Adoption and Accessibility for People with Visual and Motor Disabilities. *Proceedings of the ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*, (2009), 0-7.

32. 10. Kittur, A., Chi, E.H., and Suh, B. Crowdsourcing User Studies With Mechanical Turk. *CHI 2008*, (2008), 453-456.

33. 11. Lazar, J., Dudley-Sponaugle, A., and Greenidge, K.-D. Improving web accessibility: a study of webmaster perceptions. *Computers in Human Behavior 20*, 2 (2004), 269-288.

34. 12. Little, G. Exploring iterative and parallel human computation processes. *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems - CHI EA '10*, (2010), 4309.

35. 13. Morris, M.R., Teevan, J., and Panovich, K. What Do People Ask Their Social Networks, and Why? A Survey Study of Status Message Q&A Behavior. *Proceedings of the 2010 ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*, (2010), 1739-1748.

36. 14. Oeldorf-hirsch, A., Hecht, B., Morris, M.R., Teevan, J., and Gergle, D. To Search or to Ask : The Routing of Information Needs Between Traditional Search Engines and Social Networks. *Human Factors*, (2014), 16-27.

37. 15. Paul, S.A., Hong, L., and Chi, E.H. Is Twitter a Good Place for Asking Questions ? A Characterization Study. *Artificial Intelligence*, (2011), 578-581.

38. 16. Quinn, A.J. and Bederson, B.B. A Taxonomy of Distributed Human Computation. .

39. 17. Takagi, H., Kawanaka, S., Kobayashi, M., Itoh, T., and Asakawa, C. Social Accessibility: Achieving Accessibility through Collaborative Metadata Authoring. *ASSETS 2008*, (2008), 193-200.

40. 18. Teevan, J., Morris, M.R., and Panovich, K. Factors Affecting Response Quantity , Quality , and Speed for Questions Asked via Social Network Status Messages. *Methodology*, .

41. 19. Wu, S. and Adamic, L. Visually Impaired Users on an Online Social Network. *Proceeding of the thirty-second annual CHI conference on Human factors in computing systems - CHI '14*, (2014), 1-10.

About the Authors:

Erin Brady received her Ph.D. in Computer Science from the University of Rochester, advised by Jeffrey P. Bigham. She will be joining the Department of Human-Centered Computing in the School of Informatics and Computer at Indiana University-Purdue University Indianapolis this fall as an assistant professor.