



City Research Online

City, University of London Institutional Repository

Citation: Morrison, C., Cutrell, E., Grayson, M., Becker, E. R. B., Kladouchou, V., Pring, L., Jones, K., Faia Marques, R., Longden, C. & Sellen, A. (2021). Enabling meaningful use of AI-infused educational technologies for children with blindness: Learnings from the development and piloting of the PeopleLens curriculum. In: UNSPECIFIED . Association for Computing Machinery. ISBN 9781450383066 doi: 10.1145/3441852.3471210

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/32565/>

Link to published version: <https://doi.org/10.1145/3441852.3471210>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Enabling meaningful use of AI-infused educational technologies for children with blindness: Learnings from the development and piloting of the PeopleLens curriculum

CECILY MORRISON

Microsoft Research

EDWARD CUTRELL

Microsoft Research

MARTIN GRAYSON

Microsoft Research

ELISABETH RB BECKER

Microsoft Research

VASILIKI KLADOUCHOU

Microsoft Research

City University of London

LINDA PRING

Microsoft Research

Goldsmiths, University of London

KATHERINE JONES

University of Bristol

RITA FAIA MARQUES

Microsoft Research

CAMILLA LONGDEN

Microsoft Research

ABIGAIL SELLEN

Microsoft Research

Novel AI-infused educational technologies can give children with blindness the opportunity to explore concepts learned incidentally through vision by using alternative perceptual modalities. However, more effort is needed to support the meaningful use of such technological innovations for evaluations at scale and later wide-spread adoption. This paper presents the development and pilot evaluation of a curriculum to enable educators to support blind learners' self-exploration of social attention using the PeopleLens technology. We reflect on these learnings to present four design guidelines for creating curricula aimed to enable meaningful use. We then consider how formulations of "success" by our participants can help us think about ways of assessing efficacy in low-incidence disability groups. We conclude by arguing for our community to widen the scope of discourse around assistive technologies from design and engineering to include supporting their meaningful use.

CCS CONCEPTS • Human-centered computing • Human computer interaction (HCI) • Accessibility • Empirical studies in HCI

Additional Keywords and Phrases: Human-centred AI, user evaluation, accessibility, blindness, visual impairment, children, disability

ACM Reference Format:

First Author's Name, Initials, and Last Name, Second Author's Name, Initials, and Last Name, and Third Author's Name, Initials, and Last Name. 2018. The Title of the Paper: ACM Conference Proceedings Manuscript Submission Template: This is the subtitle of the paper, this document both explains and embodies the submission format for authors using Word. In Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY. ACM, New York, NY, USA, 10 pages. NOTE: This block will be automatically generated when manuscripts are processed after acceptance.

1 INTRODUCTION

Children with blindness cannot access the incidental learning that underpins most early learning. It is estimated that about 80% of early years learning is done through watching others [40], from figuring out the complex coordination of movements for making a sandwich to understanding the nuanced use of the body and movement for social interaction. As a result, educators of children with blindness spend significant effort habilitating a wide range of concepts and skills beyond the core academic curriculum to ensure self-determined participation in the wider community as adults [17]. Well-designed technologies can provide an alternative to prescriptive teaching approaches (e.g. social skills training [37]) by providing access to the world through alternative perceptual modalities to enable self-exploration ([24,29]).

Technological innovations that *enable* self-exploration however, require more considered effort to support meaningful use than those that provide *access* to existing curricula (e.g. accessibility of science lessons [42]). Educators need support to understand how they might productively use these technologies in meaningful ways within their educational environments; they also require evidence of their efficacy to prioritise investment and use (as done with [29]). In this paper, we discuss the development of a curriculum to support meaningful use of an AI-infused technology, exemplifying one mechanism to support evaluation at scale and consequently wide-spread adoption of novel educational technologies for children with blindness.

We ground our discussions in the example of the PeopleLens, a new AI system to enable children born blind to experience social agency and develop the range of social attention skills needed to initiate and maintain interactions [13,27]. PeopleLens uses a head-mounted augmented reality (AR) device designed to enable children with blindness to gain a dynamic, real-time understanding of their immediate social environment through spatial audio. It dynamically tracks people within four meters of the wearer, updating a world state model depending on the wearer's field of view. It identifies registered people and detects if they are looking at the wearer. In doing so, the

PeopleLens can provide an enriched understanding of social context [28,49], such as who is in the room and how they relate to the wearer. It can empower children with blindness to proactively choose their interactions whether it be to socialise, ask for help, or simply moderate their behaviour appropriately in front of teachers.

In particular, we discuss the development and piloting of an accompanying curriculum of games for the PeopleLens that can be played with peers and facilitated by educators. While social skills curriculums exist for children with blindness (e.g. [15]) and more in-depth ones for children with autism (e.g. [12,16]), these do not account for the opportunity of AI technology to support child-led exploration of social concepts at their core. Moreover, they do not reflect the new types of information potentially available to a child from an AI system nor account for the constraints of the technology. We illustrate how a curriculum can both shape meaningful use and ensure robust system performance, highlighting the important role technologists can play in contextualizing technologies for adoption.

While our community has historically limited its focus to the design and engineering of technologies, we posit that a wider scope is needed to achieve meaningful use for evaluation, and later wide-spread adoption, for many innovative AI-inspired educational technologies for children with blindness. Through the detailing of the PeopleLens curriculum and reporting of pilot study results, we illustrate how interconnected technology development is with its meaningful usage and assessment of efficacy in real-world settings. Specifically, we make the following contributions:

1. An example of a curriculum developed to account for both the new opportunities as well as the new constraints that the PeopleLens AI technology afforded.
2. Design guidelines for researchers creating curricula: 1) assessing the boundary between enablement and prescription; 2) empowering educators to adapt it; 3) fostering mutual experiences with peers; and 4) explicitly accommodating technological limitations derived from a pilot study of the curriculum with educators and learners from three locations.
3. A discussion of how we can understand and measure efficacy for AI technologies used by diverse, low-incidence disabilities grounded in participants' notions of success.
4. An argument to our community to widen the scope of discourse around assistive technologies to include approaches to support their meaningful use.

2 RELATED LITERATURE

We first review the motivation to support children born blind develop their social agency and requisite social attention skills. We then discuss learnings from related technologies and briefly overview the challenge of measuring efficacy.

2.1 Developing Social Agency and Skills

Children born blind often have substantial difficulties with social interaction, with between one- to two-thirds of these children meeting the criteria for autism spectrum conditions [34,35,47]. For example, if a blind child is not sure where their conversation partner is in space, they may have difficulty aiming their voice at their partner, or they may position their body in a way that makes them seem unengaged, such as putting their head on the table. Linguistically advanced blind children may struggle with maintaining a topic of conversation and instead talk only about something of interest to them [48]. Most noticeably, many blind children find it difficult to establish and maintain friendships with those of a similar age despite their intense desire to do so.

Despite the challenges with social interaction that many blind children have, resources and interventions are limited [52]. Manuals from teaching hubs for blind children (e.g. [15]) offer activities to improve social interaction skills, such as conversation skills and social boundaries. Social skills training has been found to be effective and generalised in peer-mediated groups for visually impaired children aged 7-12 years old [37]. Play-based approaches have also been explored to actively facilitate joint attention through off-the-shelf sound toys for children ages 7 – 9 [23] as well as bespoke toys augmented with sound for children ages 4 -13 [50]. Social skills training versus approaches to develop fundamental building blocks sit at opposite extremes for teaching children with blindness social interaction.

An alternative middle ground is to take a perceptual approach to help children learn to direct themselves towards environmental stimuli, including the dynamic elements of people [17]. Audition (hearing) can provide the sense of reference or preview needed for self-directed engagement, conveying the spatial information inherent to the organization of an environment. However, given the lower resolution of this system compared to vision, specific instruction or supported exploration is needed to develop keen auditory discrimination skills [36]. The PeopleLens is designed to render a visual scene in spatial audio in order to increase the amount of information available to a child with blindness to help them build up their adaptive strategies for engaging with the world.

2.2 Related Technologies

2.2.1 Technology Research with Blind Children

Technology research with blind children has taken many shapes. Early work focused on augmenting perception, for example making graphs auditory [8] or computing experiences haptic [32]. More recently, attention has turned to building out research prototypes that endeavor to encourage more inclusive learning with peers in mainstream education through multi-sensory robots in classrooms [25] and activities on the playground [22]. Finally, bringing these two themes together, a number of technologies have been developed for specific subjects: coding [23], computational thinking [1], geography [3], and spatial cognition [33] [11]. This short overview shows a growing literature on the use of augmented perception to enable children with blindness greater autonomy and self-direction as they develop their cognitive and social skills with peers.

Despite a remarkable level of innovation, these technological explorations have been for the most part limited in scale to a single setting with a researcher present. Torino [29] is an exception in which a curriculum was created to support independent use by teachers of the visually impaired. As a result, it was deployed to 30 educational settings (mainstream schools, special schools, and home school) for use by 75 children. This example illustrates the role a curriculum can play in supporting meaningful use at scale, inspiring the curriculum reported in this paper.

2.2.2 AI Systems for Children

Foundational work looking at how children use AI systems highlights how the use of voice assistant systems might change the knowledge seeking and structuring approach of young children [20]. More recently, this work has expanded to explore the positive ways voice assistants can support self-directed learning in 5-6 year old children as well as detail the challenges children have when navigating imperfect systems (e.g. unable to include prior context need) [19]. While the research in this space is limited, it points to the opportunity for self-exploration with AI tools by children. The challenges experienced by the children when interacting with imperfect systems directs our attention to use a curriculum as a mechanism to account for system constraints in the child's experience.

2.2.3 AI Systems for Autism

AI-enabled systems for teaching social skills in children with autism spectrum conditions are emerging. These technologies aim to achieve clinical efficacy by transforming behavioural therapies into mobile formats that can be delivered in the home or school with real-time audio-visual feedback [9,18,38,51] using AR glasses. A child usually sits opposite an adult, receiving feedback about the social situation, e.g., their eye-gaze or others' emotions. The gamified tasks are generally designed around improving social outcomes as can be measured by validated clinical measures and caregiver report. Most of the research with these systems thus far have focused on small feasibility trials, without significant success in randomized clinical trials [51].

There is growing recognition of the limitations of this "fix it" perspective that focuses on regulating stereotypical behaviours and training autistic people to follow social norms [46]. Wearable AI-systems have been developed to support social interaction in real-world settings but most have only been tested in lab settings [3]. Challenges such as guaranteeing real-time, accurate and reliable information in real-world contexts, incorporating the ethical and privacy concerns of bystanders and designing technologies that fully engage the autistic user experience (i.e. multi-sensory integration, attention challenges, social inclusion) can inhibit longitudinal study [3]. This work highlights the challenges of bringing technology to fruition that open up new ways of learning, in contrast to those that fit within clinical practice.

2.3 Measuring Efficacy

Despite significant technological innovation to support the education of children with blindness, there are still significant challenges in evaluating the efficacy of intent [6]. Quantitative measures are often key to aggregation, but their appropriateness for technologies that do not aim to compensate for an impairment but rather foster an experience or community have been questioned [6]. Inappropriate measurement choices can marginalize users' voice, allowing traditional power structures to remain (i.e., the medicalization of disability); or they can encourage the discarding of technologies that might provide value to a group of people outside the current established needs (i.e., achieving interdependence rather than independence) [30]. These challenges are particularly relevant to AI technologies that aim to augment existing capabilities, such as the PeopleLens.

3 PEOPLELENS TECHNOLOGY

PeopleLens is a real-time audio experience that is responsive to the dynamic flow and movement of social interaction as well as the movement of the user's head, differing from existing technologies such as SeeingAI [41] or Orcam [31]. Provided on a head-worn augmented reality device, it provides spatial audio information about people in the immediate vicinity via speakers located above the left and right ear. It was developed in partnership with a small number of children to create an appropriate experience to support their social interactions which is reported elsewhere [27]. Detailed description of the PeopleLens prototype has been published previously [13]. In this section we provide an overview of the features and implementation as relevant to the curriculum.

3.1 The PeopleLens Experience

The PeopleLens has the following features which are activated for particular parts of the curriculum.

Person-In-Front: This feature reads out the name of a person when the user looks at them.

- All sounds are spatialised so that the person's name is heard from the direction of that person at the time it is read out.

- If the user moves his head quickly, the notification triggers when the user’s gaze crosses the nose of a person, but the sound is spatialised according to head position at the time of rendering.

Orientation Guide: This feature provides additional sound cues to support a user’s in situ understanding of the detection of bodies or faces. These cues assist the user in orienting their body and head to interact in a socially understood way.

- Spatialised percussive “bumps” are played when a body is seen.
- Bumps are followed by a name if the person is known to the system and identifiable.
- If a person has not been identified after 0.1 seconds, a click sound is played.
- An elastic band sound can be activated to help the user shift their gaze up or down to help orient towards a face.

Gaze On Me: This feature provides a spatial sound from the direction of a person who is looking directly at the user.

3.2 Technical Details




Figure 1. PeopleLens Curriculum App. Users can add players, see game instructions and tick-off completion criteria.

3.2.1 Hardware Implementation

The PeopleLens was deployed on a modified HoloLens device (without any displays) [26] with an accompanying server box. The server box, containing two 24 GB TitanX GPUs, was used for the computation underpinning the pose and identity recognition models. The server box was provided to participants in a wheeled case to enable movement and storage. In some situations, it required 10 minutes set-up before use if it had been stored away. The external display used in [27] was not deployed to schools due to inability to certify an appropriate battery for use when a researcher is not present.

The management of the PeopleLens and the presentation of the curriculum was done with an app supplied on a tablet device (See Figure 1). This enabled teachers to train recognizable people as well as to select activities. A

selected activity in the curriculum loaded the appropriate audio experience, with opportunities for testing or doing real-world activities with different configurations of the audio experience. The app also supported data collection about the experience, including ways to delineate the user and check-off success criteria for each activity.

3.2.2 *People Tracker Implementation*

The user experience of the PeopleLens is driven by a heuristic tracking algorithm that brings together the various underlying models (e.g. pose, identity) into a persistent world state. The key purpose of the tracker is to tie together the individual observations of each model over time, accounting for the dynamic world of people who move and turn frequently. As such, the key element of the tracker is to determine whether a new observation of a person belongs to an existing person track or should spawn a new track.

There are three situational requirements for the people tracker to work well: 1) directional stability of tracked people; 2) a managed field of view; and 3) a limited number of people in the scene. Due to unreliable distance measurements calculated from 2 dimensional images that change significantly frame-to-frame, the people tracker combines new observations if they are on or close to the same directional vector from the user's head. As the PeopleLens is head-mounted and therefore moving rapidly with the head, the field-of-view, although 160 degrees, needs management. People are only timed out after 10 seconds of not being seen to account for the volatility of movement. However, this can lead to ghosts of people who are no longer there. And finally, large numbers of people significantly slow the frame rate, making predictions less accurate.

To get the best from the tracker, we aimed to encourage the following usage through the curriculum:

- Ensure that tracked people move minimally, focusing on a moving user to achieve dynamic interaction.
- Encourage people to move in front of the user where they can be tracked and not behind, causing them to be ghosts whose tracks are not updated.
- Limiting usage to 8 people maximum.

4 PEOPLELENS CURRICULUM

The PeopleLens Curriculum is a set of games that children can do with their peers in school or home settings supervised by an educator (either a teacher or parent). The games support self-directed exploration of the foundations of social attention well-supported by the PeopleLens in a systematic way. In this section, we begin by overviewing the aims of the curriculum and the development team before defining social attention and the design principles that underpin the curriculum. The final two sections describe the curriculum and elements specifically intended to support educators.

4.1 Curriculum Aims and Overview

The curriculum was developed to enable the meaningful use and validation of the PeopleLens in preparation for a long-term field study (6-9 months) facilitated by children's educators. While such studies are rare due to logistical challenges [6], they are critical to understanding whether a technology will ultimately be scalable to achieve significant impact.

In designing the curriculum, we aimed to achieve the following:

- 1) Enable the technology to fit within the educational context and align to resources an educator is accustomed to working with by providing a set of learning concepts mapped to activities that could be fruitfully explored with the PeopleLens “out-of-the-box”.
- 2) Provide a usage context that supports exploration while accounting for and engaging peers.
- 3) Ensure the robustness of the experience by embodying appropriate boundaries of use.

4.2 Development Team

A review of the literature indicated that there was no existing curriculum or intervention that was appropriate to adapt directly for use with the PeopleLens. Three of the authors worked together over a period of three months (October 2019- December 2020) to develop a set of learning objectives and linked curriculum of activities. The first curriculum team member was a professor of psychology with a wealth of research experience in blind children and social interaction; the second curriculum team member was a qualified speech and language therapist who was pursuing doctoral research; and the third curriculum team member was a human-computer-interaction researcher deeply familiar with the PeopleLens technology and who works extensively with children who are blind. The discourse of these authors aimed to bring a balance between practical skills, longer-term outcomes in social communication and respecting the agency and talents of children with blindness.

4.3 Defining Social Attention

A key step in creating the curriculum was to define social attention as relevant to children who are blind. In the spirit of self-determination [17], we wanted to support children to build skills in directing attention – their attention to others and others’ attention to them. Directing attention is core to social agency. Many children who are blind must wait for others to approach and identify themselves, giving little opportunity for these children to choose their interactions.

The term social attention was coined in the psychology literature and has been used to describe a variety of phenomena, including: 1) *behaviour* intended to coordinate attention during interaction with others; 2) *motivation* to engage with others; and 3) *attention* (movements of orienting, focusing and disengagement of the visual system) in the context of social streams of information [39]. While the most prominent use of the term *social attention* refers to nonverbal social communication [39], there are also key spatial elements to social information streams. Spatial attention limits the processing of environment stimuli to a specific location depending on how the body is positioned [7].

Spatial attention is characterized slightly differently in literature that captures perceptual approaches used to teach orientation and mobility to children with blindness [17]. Spatial information is needed for spatial referencing and preview of the environment to facilitate intentional interaction. This perceptually requires feature discrimination of a scene, event updating as the scene changes and perspective updating as the person moves through it – coming together in a dynamic spatial map. With the absence of visual focus, a child must develop other mechanisms to direct their own attention to auditory or tactile stimuli in this dynamic spatial map. No matter how achieved, we refer to this phenomenon as spatial attention. The curriculum focuses mainly on the spatial aspects of social attention.

4.4 Curriculum Design Principles

4.4.1 A Set of Structured Games

The first design principal was to create a structured set of activities to support self-exploration of social attention mechanisms. In the early development of PeopleLens, it became clear that most children we worked with had missed much of the visual incidental learning that contributes to an understanding of social interaction, e.g., people look at you because they want to interact or that one walks around groups of people not through the middle. It was not *meaningful* to provide additional social information to children if they did not know how to use it. It was observed however, that simple, structured activities allowed children to acquire this social understanding quite quickly and directed educators in how to support the learning process. As games can provide a shared social context, we choose to create a set of games closely linked to specific learning outcomes.

4.4.2 Focus on PeopleMaps

A second design principal was to focus on the spatial aspects of social attention, capitalizing on the potential of the PeopleLens technology to provide an alternative auditory spatial channel for social information. As a result, the main emphasis of the curriculum was on supporting blind children to use their heads and bodies to build and maintain a dynamic *PeopleMap* of the people around them; and then use that *PeopleMap* to effectively signal communicative intent to others in order to direct their own and others' attention in social interaction. Emphasis was placed on creating an 'image' of the social environment that enables self-directed interaction by focusing on the three articulated needs for spatial referencing and preview: feature identification (people in the *PeopleMap*), event updating (updating the *PeopleMap* as others move); and perspective updating (updating one's relationship with the *PeopleMap* as the user moves).

4.4.3 Support all Phases of Social Interaction

Our final guiding principal, drawn from the literature on social interaction, was to ensure that we supported all phases of social interaction – initiating, maintaining and switching [45]. Each of these phases requires different actions from the communicator. We also considered the role of movement in the interaction. Interactions between people seated around a table differ from having to join a group standing on the other side of the room. We addressed both of these situations in the curriculum with the assumption that younger children will begin by focusing on initiating interaction in small, structured interactions, while older children will be keen to learn how to join a small group of friends hanging out in a corner.

4.5 Curriculum Description

The PeopleLens Curriculum provides 17 games organized in two modules, available with the learning objectives listed in the supplementary materials. Module 1 focuses on interactions with a small number of people, mainly seated around a table or in a circle. Module 1 is intended to set the foundation for interaction, focusing on building and updating a *PeopleMap* in order to initiate interactions with others. The module starts with the basics of a *PeopleMap*, connecting people and location. It then progresses to monitoring change as people move around; and finally, the focus shifts to utilising the *PeopleMap* to respond to interactions. Games were drawn from published manuals used by Speech and Language therapists as well as children's games directories found online.

Module 2 is structured around a single game with multiple levels and develops interaction complexity by encouraging the child to seek out interactions, while learning to establish, maintain and shift their attention during

group interaction. The PeopleLens provides an enriched set of cues to help the child achieve more complex interactions. Once a child is able to initiate interaction by approaching someone outside their immediate interaction distance, a number of cues are used to help them locate people and find faces in order to interact. The games and additional PeopleLens cues support exploration of how the child might establish and maintain interaction with a group of two people. The final exercises focus on temporarily shifting attention away from, and then back to, a small group interaction.

To imagine these games more concretely we provide an example from each module. One of the first games in Module 1 is Maestro in which supports the child to: 1) connect people with their names and positions in space; and 2) to use head and/or body to initiate interaction with a person of preference. People sit around a table facing the child. The child assigns each person a sound. The child then turns towards people in any sequence to get a symphony of sounds. One of the final games in Module 2 aims to supports the child: 1) to prioritise an interaction, selecting what to attend to (selective attention); and prioritise an interaction based on social norms. A final level of a search and rescue game, the following instructions are given: You are busy cracking codes so that you can use your radio to call for help. You hear someone call for help that you need to rescue. On route, you hear another sound - a person, a baby, or an animal. You only have time to rescue one. Who will you choose? If you touch the enemy, nobody escapes. All games can be viewed in the Supplementary Materials.

4.6 Supporting Educators

Games were provided to educators in a printed booklet for easy reference. Each game was presented in the same manner: game objectives, game explanation, teacher tips, increasing the challenge and completion criteria. The objectives linked to the learning outcomes were intended to enable educators to appropriately adapt the game or instructions to the interests and age of the child. This was important given the wide range of children who used the same curriculum. Teacher tips and increasing the challenge helped educators understand the ways that they could support a child's exploration as well as ensure the system functioned well. Completion criteria were intended to increase the confidence of educators and help them choose the right pace to progress through the curriculum, despite variable familiarity with the adaptive strategies that children with blindness might develop.

Games, which can be found in supplementary material, were intended to be fun for all, motivating peers to participate freely. For example, "Who stole the cookie from the cookie jar?" is a game that younger children might play in their classrooms anyways. Educators were also encouraged to allow *all* students to have a turn using the PeopleLens, rotating between children after each game. This approach was intended to decrease stigma of the blind child's exploration, putting an emphasis on playing with cool technology instead. We recommended that sighted children use a blindfold to make the games meaningful and to gain perspective on auditory perception of the environment.

5 CURRICULUM PILOT STUDY

A pilot was run to explore the experiences of children (learners) and educators with the curriculum and the underlying concepts before conducting a full longitudinal study. The pilot was intended to ensure that the design choices in the curriculum were sound and respectful. There is a fine line between AI technologies enabling new capabilities and imposing a specific view on the world. We use the voices of our participants to demonstrate the potential boundaries of what enablement through curriculum could mean. In doing so, we draw on the voice of

educators and the young learners in their definition of success as a means to reflect upon how curriculum might play a role in measuring that success. The pilot study aims to address the following two questions:

R1: Did the curriculum usefully structure user exploration of social attention skills with the PeopleLens for learners, educators and peers?

R2: How did the curriculum shape educators' and learners' notions of learning development or "success" after using the PeopleLens?

5.1 Participants

Participants were recruited opportunistically from a network of educators who had previously engaged in technology research [25,29]. As this study required intensive support from educators, opportunistic recruitment enabled the researchers to be certain that educators understood the demands of using the technology and some of the challenges of prototype technology – in this case the size of the GPU box made it is effortful to move around. Three children (learners) took part, a small number chosen to reduce the demand on this low-incidence disability community [6,21].

As described in Table 1, all children were brailleists and therefore severely visually impaired as defined by the WHO definition [6]. We further collected information on the types of visual stimuli the children had at the time of the study or have had previously, including colour and shape perception, as residual vision can significantly impact a child's development. The children spanned different ages but all were boys. The study took place in either the home or school setting, depending on the child's educational setting (e.g. home school) and COVID-19 pandemic restrictions in place. All children were working at age-appropriate level as reported by their educators, but they also all had additional diagnoses as is common in this cohort. For all the children, the level of social interest, defined in the following section, was reasonably high, suggesting that difficulty engaging with others is not a matter of disinterest. Consent was gained from all learners, their educators and the parents of other children (peers) involved.

Table 1: Participant characteristics including: study ID; age; gender; level of vision--light perception, shape perception and color perception-- as defined by yes (Y), no (N) or previously (P); number of siblings; study location of home or school; additional diagnoses; and aggregated social interest as defined in 5.2.1 ranging from 0 (no social interest) to 4 (high social interest).

CHILD	AGE	GENDER	LIGHT PERCEPTION	SHAPE PERCEPTION	COLOR PERCEPTION	# OF SIBLINGS	STUDY LOCATION	ADDITIONAL DIAGNOSES	SOCIAL INTEREST
P1	13	M	Y	P	P	0	Home	ASD, dyspraxia, sensory modulation disorder	3.5/4
P2	7	M	N	N	N	1	Home	Sensory processing difficulties	3.75/4
P3	11	M	Y	N	N	2	School	Septo-optic dysplasia, epilepsy	3/4

5.2 Data Collection and Analysis

5.2.1 Social Interests Profile

The questionnaire for social interest was developed to provide a brief overview of whether social interaction challenges came from a lack of interest or skills. It contained five questions selected from the Prosocial and Peer categories of Strengths and Difficulties Questionnaire [12], the Social Relations and Interests categories of Children's Communication Checklist [4] and the Social Skills category of Social Skills Rating System [14]. Some items were rephrased to avoid misinterpretation for children with blindness. For example, the question 'Does the learner play with others?' was adjusted to 'Is the learner interested in playing with others when s/he has the chance?' In some cases, children with blindness do not have opportunity to play with others even if they want to.

5.2.2 System Log Data

A range of system log data was captured. From the app, the wearer/user was recorded along with timestamps of the start and stop of all games. Educators were encouraged to tick-off the completion criteria for each game in the app, with timestamps captured. Finally, anonymised data was collected from the PeopleLens device. This included a 3D rendering of the space and the 6DOF measures of the user's head movements. Other data, such as the head positions and locations of people identified were also post-processed. No image data was saved. The data was aggregated in a bespoke visualization software that enables various types of visualizations, including at the user activity level as well as playback of specific games as shown in Figure 2.

5.2.3 Educator and Learner Experience Questionnaires

The educator and learner questionnaires were created to understand the experience of using the PeopleLens and curriculum as well as articulate learnings that were important to those that used it. The educator questionnaire data included: 1) Likert-scale questions to understand ease-of-use and perceived usefulness; and 2) open-ended feedback on what worked well and what did not. A follow-up interview allowed deeper discussion of points raised in the questionnaire. See Appendix 2 in Supplementary Materials for documentation of questions used.

The learner questionnaire was provided via the PeopleLens App. It was intended to be read out to the learner by the educator with the answers audio recorded. We chose this mechanism to avoid challenges from children with blindness accessing written material in a variety of ways. It also allowed the educator, the person who knows the child best, to support that communication process if needed. The questions were structured to gain an understanding of how the learner experienced the PeopleLens without potentially trying to please the educator or researchers. For example, we ask the learner to describe it to a friend and whether they would like to keep it.



Figure 2: The bespoke software used to assess engagement with the PeopleLens curriculum.

6 CURRICULUM PILOT FINDINGS

The curriculum addresses R1 and R2 in subsequent sections with the first section looking at the three articulated aims of the curriculum (as stated in Section 4).

6.1 Curriculum for Meaningful Use

R1: Did the curriculum usefully structure user exploration of social attention skills with the PeopleLens for learners, educators and peers?

6.1.1 Curriculum as Learning Structure

The three learners used the PeopleLens and curriculum consistently across the time period for which they had access as shown in Table 2. Over a period of 43 – 69 days learners completed the final task in the module of choice in the curriculum. This was Module 1 for P2/3 and Module 2 for P1 whose previous experience of the PeopleLens prepared him for more advanced challenges. All educators mentioned that COVID-19 social distancing restrictions significantly limited usage, including the need to skip activities, do them with family at home, or have an overall shortened period of access to the PeopleLens. Nonetheless, curriculum progression was mainly linear with some activities skipped to meet the needs of the learner or setting restrictions.

Table 2: System usage data for participants that covers: number of distinct sessions; number of activities activated on the app; total time of usage in minutes, average length of sessions and activities in minutes; period of usage in days; total number of wearers (including educators and peers); and final activity completed as defined by the module and game number. The number in parentheses includes values for additional wearers.

PARTICIPANT	NUM SESSIONS	NUM ACTIVITIES	TOTAL USAGE (MIN)	AVERAGE SESSION LENGTH (MIN)	AVERAGE ACTIVITY LENGTH (MIN)	DURATION OF USE PERIOD	NUM WEARERS	FINAL ACTIVITY
-------------	--------------	----------------	-------------------	------------------------------	-------------------------------	------------------------	-------------	----------------

P1	8(9)	24(29)	113(122)	14:15(13:30)	4:45 (4:15)	69 days	3	M2:9
P2	15 (18)	39 (58)	90 (110)	6 (6)	2:15	55 days	2	M1:6
P3	10 (10)	20 (32)	67 (106)	10:40 (6:45)	3:20	43 days	6	M1:7

Activities happened in short bursts and lengthened as the participant’s age increased, ranging from 2:15 – 4:45 minutes. Sessions included a series of activities which lasted between 6 – 14:15 minutes, with pauses in between for a variety of reasons: sharing the system with other children (P3), waiting for a natural moment of need to occur (e.g. when someone entered the room) (P2), or just discussing the experience and learnings (P1). This suggests that activities were a starting point for learning, but much work was being done by educators to contextualize the learning for the learner. Educators commented that the short bursts made it easy to fit this non-curricular engagement into the day and include peers in the school setting as appropriate.

All educators adapted the curriculum to meet the needs and age of the participating learner. In Game 2: The Maestro, people sit around a table facing the learner. The learner assigns each person a sound. The learner then turns towards people in any sequence to get a symphony of sounds. The aim of the game is to encourage the learner to connect people with their names and positions in space; and use their head and/or body to initiate interaction with a person of preference. In the case of P2, sounds were turned into motivating songs and the learner was encouraged to stay facing the person with the MP3 player to hear their whole song. Songs were chosen to represent the person holding the MP3 player building on P2’s social engagement with people by singing “their” song.

In contrast for P3, who was substantially older and using the PeopleLens in a school setting, the game unfolded in a different way. The first attempt at the Maestro game was unstructured, with everyone playing a percussion instrument when looked at by the learner. The game then evolved into a deliberate and purposeful process with each person playing a different note on a xylophone. In the final playing of the game, P3 tried to play a tune with xylophone notes. The group also experimented with each person playing their xylophone note only once as well as continually until P3 looked away. In addition to the skills that P3 gained in orienting towards people he liked experiencing social agency by being the conductor in charge of the group.

Educators expressed no difficulty or lack of confidence in adapting the curriculum, as illustrated in the vignettes above. Additionally, the response to the Likert-scale question found in Table 3 ‘I found the Scheme of Work easy to follow as a teacher’ was positive (4.67/5). Educators were confident enough to go “off-curriculum.” The educator for P2 used the PeopleLens during teachable moments, e.g. when someone just entered the house, in an attempt to stimulate an interest in engaging with people outside the immediate social zone. In the case of P3, the educator often let the structured activity lapse into freeform conversation. As researchers, we might consider this the ideal scenario in which the children are positioned for easy recognition by the PeopleLens but the interaction can be natural.

Overall, educators were confident in the learnings of their students. The Likert-scale questions showed high ratings of the PeopleLens as a tool for teaching children with blindness about social interaction (5/5) with a strong desire to use it again (5/5). These sentiments were expressed in the educators’ articulations of student learnings. For example, P3’s educator wrote: “He definitely developed the idea that when somebody is talking to him or he is talking to somebody else he needs to face them and look up” [to direct their attention to him]. The above findings suggest that the curriculum played an important role in facilitating exploration and understanding of the spatial concepts of social attention and that educators were able to use and adapt appropriately “out-of-the-box.”

The biggest challenge with the technology during this pilot was the bulkiness and weight of the server, curtailing exploration that educators had in mind to enable. This, along with some difficult sequencing required when starting

all the components of the system, likely accounts for the Likert-scale score on the difficulty of using PeopleLens (2.33/1 (ideal response)).

Table 3: The responses for each Likert-scale question as well as the aggregate.

Question	P1	P2	P3
I think PeopleLens tool is useful for teaching social interaction skills to blind children.	5	5	5
I do not think the PeopleLens tool helps communication partners (other people) interact with the blind child.	2	1	2
I think the PeopleLens tool helps blind children interact more easily with their peers	5	3	5
I found the PeopleLens tool difficult to use as a teacher.	2	2	3
I found the Scheme of Work easy to follow as a teacher.	5	5	4
I would like to use the PeopleLens tool with other learners in the future.	5	5	5

6.1.2 Curriculum as Mutual Experiences

PeopleLens was intended to support structured self-exploration of social concepts. The data suggests that learners were able to experience the empowerment and enjoyment of joining in with social norms. P1 / P2 said:

“I just seem to never want to take it off. When it’s on my head, even though it’s very heavy, I like to keep it on my head. It just gives me so much information that I would otherwise not have had that people around me have all the time. It gives me more freedom. It gives me more possibilities, like I can just say hi to people.” (P1)

“It was so much fun. I liked it when my friends came together, we chatted about it and we got to socialise.” (P2)

Learners did not mention or allude to any sense of dissatisfaction or enforced social behaviour. There was no referencing by learners of the constant reminders to keep their head up or look at another child that the educators spoke about. As articulated by the educator of P1:

“[P1] seemed to really get a feel for the reactions he could get from others from how he placed his body and especially angled his body and face.”

This suggests that the PeopleLens was a good mix of structure and focus for educators, without inhibiting the learner’s experience of social agency.

The curriculum was designed to be used with peers. While it was only used in one school setting due to COVID-19, the games seemed to support learning. In the school setting, the participating peers all used the PeopleLens during each session. The educator said that the children enjoyed doing something non-academic together for 10-15 minutes as well as liked the novelty of the experience. This is reflected by the scores on the Likert-scale questions around enabling peer interaction (4.33/5). The framing of the PeopleLens as a benefit to all seems to have worked well in enabling a mutually positive experience between the peers and the learner, avoiding potential stigmatization from using the technology [43,44].

Currently the curriculum learning outcomes are mainly focused toward the child with blindness. While this generally seemed to fit with the expectations of educators, potentially more could be done to move the benefit of

the experience for peers from novelty to being a more empathetic social communicator. This may be reflected in the Likert-scale question that captured the role of PeopleLens in helping communication partners (1.67/5). There is room here to consider the learnings, and not just participation, of peers and how the curriculum and technology facilitate mutual engagement. Nonetheless, the data suggests that the curriculum succeeded in enabling exploration while accounting for and engaging peers.

6.1.3 Curriculum as Technology Boundary

The curriculum was specifically designed to minimize the weaknesses of the PeopleLens technology within the scope of usage for benefit. For example, games in Module 1 were set up in a circle or around a table to reduce the likelihood of movement or needing to recognize the side of someone's face. Looking at the log-data, we can see that extra tracks were spawned and appear as extra bump sounds (not names) in the experience. There are nearly 3 times fewer extra sounds in Module 1 than in Module 2 when people are moving (average of 2.5/m versus 6/m). Digging deeper into the data, we can see that the extra track variation is very high with most games having none, and a few, often played on the same day, having large numbers. We surmise from the 3D reconstruction data that some setups do not work well for recognition. While curriculum design choices certainly help, further data analysis proves that even more could be done to aide recognition.

We also saw examples in which the curriculum did not manage to successfully bound technology usage. The Shuffling Tornado game was set up so that peers stand spaced out in a semi-circle in front of the learner and then change places. The learner then has to guess the new order. This game was played by P2 at a table against a wall. As a result, the shuffling took place behind the learner rather than in front. The tracker was unable to keep a view on where people were. This resulted in the position of a person track not updating and the wrong name being read out. The educator described this as a confusing situation for the learner, who could smell the person who was actually there. While the game was crafted to avoid this known difficulty with the person tracking algorithm, the on-the-fly adaptations in a family home still led to non-performance on several occasions.

AI technologies can be unpredictable and that is even more true of novel technologies which are pushing technical boundaries as is the case of PeopleLens. Yet, deployment is critical to improve the underlying technology and understand how it is most likely to enable intended users. These findings demonstrate that a curriculum can be an effective way of providing appropriate bounds on a new technology to ensure its robustness in long-term evaluations without the researcher being present.

6.2 Curriculum as Shared Understanding of Success

R2: How did the curriculum shape educators' and learners' notions of learning development or "success" after using the PeopleLens?

6.2.1 Success as skills

The curriculum provided a shared context for educators to talk about the success of their learners with researchers. Educators were articulate and specific about the concepts that they wanted to work on with their learners: attending to a person talking (P2/P3) and joining a group of peers (P1). The curriculum played an important bridging role in helping educators understand the scope of potential success. It provided a context for speaking about the learning outcomes and skills related to the goals for the learner. For example, P3 chose the learning

outcome of looking at others in discussion with his educator as part of a larger goal-setting exercise at the beginning of term.

6.2.2 Success as discovered achievements

Educators' articulations of success, however, went beyond the curriculum. Despite the curriculum as the center piece of discussion between educators and researchers, written documentation of their experience and post-interviews illustrated more nuance to how success emerged during the pilot. Two examples include an increase in verbal fluency when attention is given (P2) and an increased ability to switch attention between people and meet their eye gaze by looking up or down a face (P1). The educators gave the impression that these were discovered achievements rather than intended learning outcomes.

6.2.3 Success as exposure to new ideas

Perhaps even more nebulous, but arguably a more powerful success, came in learnings from understanding the potential shape of social communication. Educators noted that elements of the experience that gave access to information that learners were already accustomed to, but had difficulty acquiring, such as people's names, showed immediate benefit. All educators contrasted this type of learning with helping a learner understand elements of social communication of which they had no experience, e.g., that someone looking at them could indicate a desire for interaction. This sense of opening up a new point-of-view for the learner, was captured by one educator as: "For [P2], [the PeopleLens experience] inserted the idea that people exist in space."

In this latter example, success looked like the fundamental realization of a new idea, a step that is critical to real change in social attention. It is perhaps what we might wish most from the PeopleLens experience, a way to deeply explore social attention rather than impose specific behaviours, an approach taken with AR technologies developed to support social communication for children with autism spectrum conditions [9]. However, educators noted that in this study it took concerted practice to help learners account and respond to these new types of information. Awareness and skill were different phases of the learning process, with the generalized skill not being realized in the timeframe of the pilot.

6.2.4 Success as holistic change

Each of the educators involved wrote first of the successes that aligned to learning outcomes or specific skills (e.g., meeting people's gaze). However, all of them, also spoke more holistically about improvements. For example, in the case of P1:

As [P1]'s parent I find myself feeling quite emotional watching him when he is using the Peoplens. I don't know how much this is just about my ingrained social expectations as a sighted person but it does seem as if [P1]'s whole demeanor is more natural when using the PeopleLens. This coupled with [P1]'s great enjoyment of the experience makes me feel that it is giving him an experience that is useful and positive for him as a blind person too.

It seems that these articulations of success capture the integration of different skills, or perhaps the development of a fundamental building block, to reach a qualitative difference in how social attention is achieved.

7 DISCUSSION

Novel AI-infused educational technologies can give children with blindness the opportunity to explore concepts learned incidentally through vision by using alternative perceptual modalities. However, more effort is needed to support the meaningful use of such technological innovations for evaluations at scale and later wide-spread adoptions. This paper presents the development and pilot evaluation of a curriculum to enable educators to support learners with blindness explore social attention using the PeopleLens technology. In the discussion we integrate our learnings in designing a curriculum by offering four design guidelines for other researchers. We also reflect upon how we might assess the efficacy of such technologies given the voices of our educator and learner participants. Following a section on limitations, we close the discussion with an argument for our community to widen the scope of discourse around assistive technologies to include approaches to support their meaningful use.

7.1 Enabling Meaningful Use

Innovative educational technologies need to be used meaningfully. It may not be enough to provide learners with information in a new modality if they do not know how to use that information. As demonstrated with the PeopleLens, it was not enough to provide the names and locations of people as children with blindness have not had the early years experience with this information and could not use it effectively to manage social attention. Moreover, the detailed thought that grounds a technology design and epistemic approach needs to be communicated to those managing the technology, usually educators. This is particularly challenging when working with children with sensory impairments, as they are low-incidence and therefore geographically dispersed. The curriculum presented in this paper is one example of how meaningful use of a novel technology can be supported when a researcher does not have direct interaction with participants. In the following paragraphs, we elucidate four design guidelines for other researchers.

7.1.1 *Assess the Balance between Enablement and Prescription*

Structure is essential to support productive exploration by learners as well as guide the educator in finding appropriate usage in their educational environment. Structure does not have to equate with prescription. Although prescriptive approaches to education are common, they can disempower learners with blindness [17]. As we saw in the pilot reported here, the curriculum achieved a balance between enablement (with learners feeling empowered) while educators had the structure they required for use in their diverse contexts.

7.1.2 *Empower Educators to Adapt*

Children with low-incidence disabilities are very diverse [5]. Technologies may be used across a wide age-spectrum; contexts also vary widely from special schools and home school to bespoke mainstream inclusion approaches. Educators know the needs of their learners and are best placed to provide the experience most relevant to the development of that learner. Explicit encouragement is a good starting point to engage educators in adapting the curriculum. We also found that learning outcomes and completion criteria, as they communicated the focus of the learning, played a key role in increasing the confidence of educators in their adaptations.

7.1.3 *Create Mutual Experiences*

Assistive technologies, even enabling ones, can cause stigmatization [43,44]. Creating mutual experiences that include peers in a productive way can proactively shape views on the most productive use of a technology in

supporting the learning process. While this curriculum was successful in creating a non-stigmatized learning environment, we realize that we could have gone further in making the experience of peers not just fun and novel, but also more educational as social communicators to people with different world perspectives.

7.1.4 Accommodate Technology Limitations

AI technologies are difficult to design for as their outcomes can be unpredictable. While AI technologies have been shown to power productive educational experiences for children, their limitations can significantly dampen the experience. The PeopleLens curriculum was designed to accommodate the known limitations of the technology experience (as detailed in section 3). For the most part, this was successful with the only “catastrophic” system failure being a learner getting confused due to an adaptation of the curriculum.

7.2 Assessing Efficacy

Despite significant technological innovation to support the education of children with blindness, there are still significant challenges in assessing efficacy [6]. Many of these challenges are related to a highly diverse, low-incidence population, making it difficult to use traditional quantitative measures. Another challenge comes from the conceptualization of the technology itself. If the aim of a technology is to help learners explore and build out new, unforeseen capabilities, notions of efficacy must also be emergent. Nonetheless, uptake and wide-spread adoption of novel educational technologies depends on the ability to measure their efficacy. Below are reflections on how the educators’ own articulations of success might shape research notions of efficacy.

7.2.1 Curriculum as Measurable Skills

A skills-based approach in which a set of skills is chosen and worked on (a typical educational approach) using the curriculum can provide a way to realistically measure skill acquisition. In the case of the PeopleLens, an incidental view of this learning might be seen through logging measures such as completion criteria or using algorithms to extract demonstrations of the skills from post-processing of captured data (e.g., 6DOF measures of head position); however, a more realistic approach may be to observe the changes in these skills through quantified coding of video of a real-world activity. This skills-based approach would be one way to underpin proposals to use single person (pre/post) study designs to address the heterogeneity of learners and the benefit they receive.

7.2.2 Curriculum as Shared Understanding

Yet, the reflections of educators suggest that much of their view of “success” was outside the realms of specific skills captured in the curriculum. In some cases, success was emergent (e.g., better sustained conversation); in other cases, it was the realization of an idea (e.g., a person looks at another to direct their attention); it was also described as the integration of a number of skills to show consequential change in social communication. The learning objectives of the curriculum did, however, provided a shared understanding of what successful learning explorations might look like. Unlike the large-scale study done with children with blindness or low vision for a physical coding language [29], this pilot showed that educators had little difficulty precisely articulating what success they saw, making their communication of what learners achieved more consistent.

Using the curriculum as a way to build shared understanding with the people who will not only deliver the educational experience, but also assess it, provides a compelling alternative to typical quantitative measures of success that use standardized questionnaires. Often developed in clinical contexts, these questionnaires, such as

the Children’s Communication Checklist (CCC) [4] are difficult to interpret for those with low-incidence disabilities. Moreover, they also do not speak to the formulation of technology as an exploratory and empowering tool. Indeed, questions can feel invasive and negative, such as asking a child if they have a friend, a painful question for a child with blindness struggling to find connection. Learning objectives offer a starting point for educators to articulate and track change as well as shape the collaborative learning journey with the student, making participation in research a positive experience.

7.2.3 *Assessing Emerging Capabilities*

It is clear that ways of consistently capturing an educator’s view of change will be important, as this pilot has shown that notions of “success” change over time. Understanding a learner’s sense of change may also be important, although in this pilot we did not see the levels of self-reflection needed to rely on this alone. This may be the result of the age of the learners, or the need for more time for the individual skills to aggregate into changed social interactions. Capturing the view of participant educators and learners in a pilot gives an opportunity to identify measurable elements that contribute to efficacy. These could include changes in fundamental capabilities (e.g., dynamic spatial processing) or in higher-level aggregate skills (e.g., verbal fluency [10]). That said, there is still more research effort needed to take such an approach as the research about fundamental capabilities is often limited by existing technologies and may take long-term studies to see change.

We welcome further research into tools that might support qualitative and quantitative enquiry that can help elucidate and measure the role of AI systems in supporting the development of children with disabilities.

7.3 Study Limitations

This study focuses specifically on the experiences of educators and learners on piloting the PeopleLens curriculum. It does not frame that discussion within actual measurement of efficacy of the technology (i.e., learning outcomes achieved). Indeed, we specifically wanted to understand what success might look like before codifying it in quantitative measures. Moreover, in respect of the research burden in low-incidence disability communities [21], we only recruited three participants, the minimum number in which we felt we could get a good sense of the variation. A larger number of children and a more complete set of data is needed to validate the full PeopleLens experience and measure its impact. We hope this will become possible once COVID-19 pandemic restrictions lessen.

8 CONCLUSION

There has been significant innovation in educational technologies for children with disabilities, but evaluations of such technologies have generally been small-scale ([29] is an exception) due to challenges in supporting meaningful use [6,21] and the fragility of prototypes, particularly AI technologies [3]. We discussed the successful role of the curriculum in enabling educators across a range of settings to use a novel AI technology without confining the experience of the learners. We also illustrated the important role of the curriculum in minimizing catastrophic technology failure, a necessary requirement for long-term deployments of prototype technology. Thinking about how we support meaningful use of the technologies we create will help address the challenge of small sample size in studies with children with disabilities [5] and blindness more widely [6]. Usage can be more geographically distributed with higher confidence that the technology will be robust.

Our community also has a rising discourse about using technology to challenge the status quo, empowering people with disabilities to more actively shape the potential of technologies [2,53]. This will be difficult to achieve

if technologies are slotted into an existing eco-system that supports the model of medicalization of disability. For example, technology used in therapies for autism (e.g. [9,51]) is confined by those existing therapies and measures of success, and, if effective, might replace a skilled practitioner to drive cost down. While such an approach makes it easier to imagine and scale a technology, it limits the potential for innovation that enables people with disabilities in new ways. Broadening our scope to include a concern about scaling meaningful use is one step needed to encourage technologies that challenge the status quo.

ACKNOWLEDGMENTS

We would like to acknowledge our participants and their support networks as well as Anja Thieme for contributing to the formulation of the paper topic.

REFERENCES

1. Lúcia Abreu, Ana Cristina Pires, and Tiago Guerreiro. 2020. TACTOPI: a Playful Approach to Promote Computational Thinking for Visually Impaired Children. In *The Proceedings of the 2020 SIGACCESS Conference on Computers and Accessibility*, 1–3. <https://doi.org/10.1145/3373625.3418003>
2. Cynthia L. Bennett, Erin Brady, and Stacy M. Branham. 2018. Interdependence as a frame for assistive technology research and design. In *Proceedings of the 2018 SIGACCESS Conference on Computers and Accessibility*, 161–173.
3. Esma Mansouri Bessassi, Juan Carlos Gomez, Louanne E. Boyd, Gillian R. Hayes, and Juan Ye. 2018. Wearable Assistive Technologies for Autism: Opportunities and Challenges. *IEEE Pervasive Computing* 17, 2: 11–21. <https://doi.org/10.1109/MPRV.2018.022511239>
4. D. V. Bishop. 2003. *The children's communication checklist: CCC-2*. Harcourt Assessment.
5. Emeline Brulé, Oussama Metatla, Katta Spiel, Ahmed Kharrufa, and Charlotte Robinson. 2019. Evaluating Technologies with and for Disabled Children. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–6.
6. Emeline Brulé, Brianna J. Tomlinson, Oussama Metatla, Christophe Jouffrais, and Marcos Serrano. 2020. Review of Quantitative Empirical Evaluations of Technology for People with Visual Impairments. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–14.
7. Kyle R. Cave. 2013. *Spatial Attention*. Oxford University Press.
8. Robert F. Cohen, Arthur Meacham, and Joelle Skaff. 2006. Teaching graphs to visually impaired students using an active auditory interface. *SIGCSE Bulletin* 38, 1: 279–282.
9. Jena Daniels, Jessey N. Schwartz, Catalin Voss, Nick Haber, Azar Fazel, Aaron Kline, Peter Washington, Carl Feinstein, Terry Winograd, and Dennis P. Wall. 2018. Exploratory study examining the at-home feasibility of a wearable tool for social-affective learning in children with autism. *npj Digital Medicine* 1, 1: 1–10.
10. P. Ekman and W. V. Friesen. 1969. The repertoire of nonverbal behavior: Categories, origins, usage, and coding. *semiotica* 1, 1: 49–98.
11. Euan Freeman, Graham Wilson, Stephen Brewster, Gabriel Baud-Bovy, Charlotte Magnusson, and Hector Caltenco. 2017. Audible beacons and wearables in schools: Helping young visually impaired children play and move independently. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 4146–4157.
12. R. Goodman. 1997. The Strengths and Difficulties Questionnaire: a research note. *Journal of child psychology and psychiatry* 38, 5: 581–586.
13. Martin Grayson, Anja Thieme, Rita Marques, Daniela Massiceti, Ed Cutrell, and Cecily Morrison. 2020. A dynamic AI system for extending the capabilities of blind people. In *Conference on Human Factors in Computing Systems - Proceedings*. <https://doi.org/10.1145/3334480.3383142>
14. F. M. Gresham and S. N. Elliott. 1990. Social skills rating system: Manual. *American Guidance Service*.

15. Linda Hagood. 2008. *Better Together: Building relationships with people who have visual impairment and autism spectrum disorder (or atypical social development)*.
16. Neha U. Keshav, Arshya Vahabzadeh, Rafiq Abdus-Sabur, Krystal Huey, Joseph P. Salisbury, Rungpeng Li, and Ned Sahin. 2018. Longitudinal socio-emotional learning intervention for autism via smartglasses: Qualitative school teacher descriptions of practicality, usability, and efficacy in general and special education classroom settings. *Education Sciences* 8, 3: 107.
17. Daniel Kish and Jo Hook. 2016. *Echolocation and FlashSonar*.
18. Rungpeng Liu, Joseph P. Salisbury, Arshya Vahabzadeh, and Ned T. Sahin. 2017. Feasibility of an Autism-Focused Augmented Reality Smartglasses System for Social Communication and Behavioral Coaching. *Frontiers in Pediatrics* 5: 145.
19. Silvia B. Lovato, Anne Marie Piper, and Ellen A. Wartella. 2019. "Hey Google, do unicorns exist?": Conversational agents as a path to answers to children's questions. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children, IDC 2019*, 301–313.
20. Silvia Lovato and Anne Marie Piper. 2015. "Siri, is this you?": Understanding Young Children's Interactions with Voice Input Systems. In *Proceedings of the 14th International Conference on Interaction Design and Children*, 335–338.
21. K Mack, E McDonnell, D Jain, LL Wang, J Froehlich, and L Findlander. 2021. What Do We Mean by "Accessibility Research"? A Literature Survey of Accessibility Papers in CHI and ASSETS from 1994 to 2019. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
22. Caltenco H. Magnusson C., Hedvall PO. 2018. Co-designing together with Persons with Visual Impairments. In *Mobility of Visually Impaired People*, Pissaloux E. and Velazquez R (eds.). Springer.
23. Mary Christine Ross. *Promoting Joint Attention in Children with Visual Impairment: Proposing an Intervention Using Modified Strategies from Joint Attention Symbolic Play Engagement Regulation*. Doctoral dissertation, The Ohio State University
24. Oussama Metatla, Sandra Bardot, Clare Cullen, Marcos Serrano, and Christophe Jouffrais. 2020. Robots for Inclusive Play: Co-designing an Educational Game With Visually Impaired and Sighted Children. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–10.
25. Oussama Metatla, Sandra Bardot, Clare Cullen, Marcos Serrano, and Christophe Jouffrais. 2020. Robots for Inclusive Play: Co-designing an Educational Game with Visually Impaired and sighted Children. In *Proceedings of the 2020 Conference on Human Factors in Computing Systems* 1–13.
26. Microsoft Hololens. <https://www.microsoft.com/en-us/hololens>.
27. Cecily Morrison, Ed Cutrell, Martin Grayson, Anja Thieme, Alex S Taylor, Geert Roumen, Camilla Longden, Rita Marques, Abigail Sellen, and Sebastian Tschitschek. 2021. Social Sensemaking with AI: Designing an Open-ended AI experience with a Blind Child. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–12.
28. Cecily Morrison, Edward Cutrell, Anupama Dhareshwar, Kevin Doherty, Anja Thieme, and Alex Taylor. 2017. Imagining Artificial Intelligence Applications with People with Visual Disabilities using Tactile Ideation , pp. 81-90. ACM, 2017. In *Proceedings of the 2017 SIGACCESS Conference on Computers and Accessibility*, 81.
29. Cecily Morrison, Nicolas Villar, Alex Hadwen-Bennett, Tim Regan, Daniel Cletheroe, Anja Thieme, and Sue Sentance. 2019. Physical Programming for Blind and Low Vision Children at Scale. *Human-Computer Interaction*.
30. I. Neto, H. Nicolau, and A. Paiva. Fostering Inclusive Activities in Mixed-visual Abilities Classrooms using Social Robots. In *Companion of the 2021 Conference on Human-Robot Interaction*.
31. Orcam My Eye. <https://www.orcam.com/en/>.
32. Saija Patomäki, Roope Raisamo, Virpi Pasto Jouni Salo, and Arto Hippula. 2004. Experiences on haptic interfaces for visually impaired young children. In *Proceedings of the ICMI Conference on Multimodal interfaces*, 281–288.
33. Lope Ben Porquis, Sara Finocchietti, Giorgio Zini, Giulia Cappagli, Monica Gori, and Gabriel Baud-Bovy. 2018. ABBI: A wearable device for improving spatial cognition in visually-impaired children. In *Proceedings of the BioCas 2017 Conference on Biomedical Circuits and*

- Systems C*, 1–4.
34. Linda Pring (Ed.). 2005. *Autism and blindness: Research and reflections*.
 35. Linda Pring. 2008. Psychological characteristics of children with visual impairments: learning, memory and imagery. *British Journal of Visual Impairment* 26, 2: 159–169. <https://doi.org/10.1177/0264619607088279>
 36. Susanne Smith Roley, Erna I. Blanche, and Roseann C. Schaaf. 2001. *Understanding the nature of sensory integration with diverse populations*. Pro-Ed.
 37. Sharon Sacks and Robert Gaylord-Ross. 1989. Peer-mediated and teacher-directed social skills training for visually impaired students. *Behavior Therapy* 20, 4: 619–640.
 38. Ned T. Sahin, Rafiq Abdus-Sabur, Neha U. Keshav, Runpeng Liu, Joseph P. Salisbury, and Arshya Vahabzadeh. 2018. Case Study of a Digital Augmented Reality Intervention for Autism in School Classrooms: Associated With Improved Social Communication, Cognition, and Motivation via Educator and Parent Assessment. *Frontiers in Education* 3: 57.
 39. Brenda Salley and John Colombo. 2016. Conceptualizing Social Attention in Developmental Research. *Social Development* 25, 4: 687–703.
 40. A Salt and N Dale. 2019. *Developmental Journal for Babies and Young Children with Visual Impairment (DJVI) for Professional Use*. Retrieved from https://xip.uclb.com/i/healthcare_tools/DJVI_professional.html
 41. SeeingAI. <https://www.microsoft.com/en-us/ai/seeing-ai>.
 42. Lei Shi, Idan Zelzer, Catherine Feng, and Shiri Azenkot. 2016. Tickers and Talker: An Accessible Labeling Toolkit for 3D Printed Models. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*: 4896–4907.
 43. Kristen Shinohara and Jacob O. Wobbrock. 2011. In the shadow of misperception. *Proceedings of the 2011 CHI Conference on Human Factors in Computing Systems*: 705.
 44. S Söderström and B Ytterhus. 2010. The use and nonuse of assistive technologies from the world of information and communication technology by visually impaired young people: a walk on the tightrope of peer inclusion. *Disability & Society* 25, 3: 303–315.
 45. McKay Moore Sohlberg and Catherine A. Mateer. 2006. Improving Attention and Managing Attentional Problems. *Annals of the New York Academy of Sciences* 931, 1: 359–375.
 46. Katta Spiel, Christopher Frauenberger, Eva Hornecker, and Geraldine Fitzpatrick. 2017. When empathy is not enough: Assessing the experiences of autistic children with technologies. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2853–2864.
 47. Valerie Tadic, Linda Pring, and Naomi Dale. 2009. Attentional processes in young children with congenital visual impairment. *British Journal of Developmental Psychology* 27: 311–330.
 48. Valerie Tadić, Linda Pring, and Naomi Dale. 2010. Are language and social communication intact in children with congenital visual impairment at school age? *Journal of Child Psychology and Psychiatry* 51, 6: 696–705.
 49. Anja Thieme, Cynthia L. Bennett, Cecily Morrison, Edward Cutrell, and Alex S. Taylor. 2018. "I can do everything but see!--How People with Vision Impairments Negotiate their Abilities in Social Contexts. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 203.
 50. Suzanne H Verver, Mathijs PJ Vervloed, and Bert Steenbergen. 2019. The use of augmented toys to facilitate play in school-aged children with visual impairments. *Research in developmental disabilities* 85: 70–81.
 51. Daniels J Voss C, Schwartz J. 2019. Effect of Wearable Digital Intervention for Improving Socialization in Children With Autism Spectrum Disorder: A Randomized Clinical Trial. *AMA Pediatr* 173, 5: 446–454.
 52. E Yesilkaya, B Byrne, and G Marshall. 2019. *The Barriers and Facilitators to the Implementation of Interventions for Children with Visual Impairments, Their Parents/Guardians or Educators: A Systematic Scoping Review*. Routledge.
 53. Anon Ymous, Katta Spiel, Os Keyes, Rua M. Williams, Judith Good, Eva Hornecker, and Cynthia L. Bennett. 2020. "I am just terrified of my

future”—Epistemic Violence in Disability Related Technology Research. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–16.