Speak Up: A Case Study in the Interplay Between Culture and Technology

Amal Nanavati
Carnegie Mellon University
Pittsburgh, USA
amaln@cmu.edu

M. Bernardine Dias
Carnegie Mellon University
Pittsburgh, USA
mbdias@ri.cmu.edu

Aaron Steinfeld
Carnegie Mellon University
Pittsburgh, USA
steinfeld@cmu.edu

ABSTRACT
The ability to communicate is crucial to leading an independent life. Unfortunately, hearing impaired individuals in developing communities tend to have difficulty communicating, due to the lack of acceptance of sign language, a scarcity of speech therapists, and cultural stigmas surrounding deafness. In 2015, research interns from our team developed a suite of voice-powered games to allow hearing impaired students in a school for the deaf in India to visualize, understand, and explore their voices. Two years later, we returned to the community to conduct fieldwork and iterate on the games. This paper presents insights from that fieldwork. It details ethnographic insights about the interconnection between the community's cultural background and the games, enhancements we made to the games based on those insights, and an exploration of the potential for technological tools to build capacity in underserved communities.

INTRODUCTION
Especially in the developing world, individuals with hearing loss tend to lack the technological, societal, and educational support necessary to lead independent lives. Although the "majority of people with disabling hearing loss live in low- and middle-income countries", only 3% of developing countries' hearing aid needs are met [6]. In India, despite the Person with Disabilities Act of 1995 that makes education of disabled a government priority [27], only about 0.5% of deaf children attend a school specifically designed to meet their needs [31]. The Deaf community in India also faces a great deal of cultural stigma, including the idea that deafness implies dumbness, or an inability to communicate. Further, speech therapists in India are scarce [33] and expensive. Lastly, sign language is not a culturally accepted form of communication, due to a lack of awareness of and regional variations in the language [31, 32]. Therefore, deaf children do not have the resources or opportunities to learn verbal communication, although it is a prerequisite for integrating into mainstream society.

In 2015, four research interns from Carnegie Mellon University's TechBridgeWorld, Nanavati included, went to the Mathru Center for the Deaf and Differently-Abled (Mathru Center) for nine-weeks to investigate the potential for and develop technological tools to support the school community. One of the tools we developed was *Speak Up! Voice Powered Game Suite*, a collection of computer games intended to help hearing-impaired students explore and understand their voices. Two years later, I (Nanavati) returned to do follow-up fieldwork and iteration on *Speak Up*. I took on dual roles during this research: 1) that of a technology developer, further enhancing *Speak Up* based on community feedback; and 2) that of an ethnographer, observing the ways in which technology shapes and is shaped by the community's dynamics, and developing strategies to incorporate those insights into the technology development process. Dias and Steinfeld, faculty advisers, worked with me before, during, and after the fieldwork to provide methodological, technology development, and other relevant advice.

This paper details insights primarily from the most recent two weeks of fieldwork. It begins by describing the backdrop and methodologies for the fieldwork, proceeds to detail social dynamics and human factors in the Mathru Center community and how the games shape those, goes on to describe enhancements I made to *Speak Up* based on the aforementioned observations, and ends with a discussion of methodological insights and the potential for technology to be used for capacity building. Its chief contributions include:

1. An exploration of the complex relations between cultural factors and technology usage, and how to incorporate that knowledge into the technology development process;
2. Initial findings about the potential for embodied physical agents to enhance educational technologies for developing communities;
3. An explicit recognition of the role(s) technology has in building capacity within the community and a case study in how to integrate that into the technology development process.

BACKGROUND
This section gives necessary background about the Mathru Educational Trust for the Blind, the environmental, social, and technical contexts of the Mathru Center, the *Speak Up! Voice Powered Game Suite*, and past research that is related to our work.

Institutional Partner
The Mathru Educational Trust for the Blind[1] was founded

[1] http://www.mathrublindschool.org/

by Ms. Gubbi Muktha. After a vehicle accident that left her debilitated, Ms. Muktha began frequenting a disability and rehabilitation center, where she saw first-hand the plight of many blind children in Bangalore. This inspired Ms. Muktha to start the Mathru School for the Blind, a free residential school for blind students, in 2001. In 2012, the Trust expanded its focus by founding the Mathru Center for the Deaf and Differently-Abled, another residential campus that includes a school for the deaf, a classroom for students with multi-sensory impairments, and a classroom for students with multiple disabilities. The Trust also runs an educational outreach program for individuals who are deaf and blind, computer and vocational training for local women, and other charitable projects [20, 13, 14].

Researchers from Carnegie Mellon University's TechBridgeWorld have been working with the Mathru Educational Trust for the Blind since 2006. Select technologies that have resulted from this partnership include the Stand-Alone Braille Tutor [3] and a Tactile Graphics converter [5] for the School for the Blind. In 2013 and 2015, research interns from TechBridgeWorld conducted a needs assessment at the Mathru Center, and found that: 1) it was difficult for new teachers to learn the job due to a lack of teacher training materials, 2) it was difficult and time-intensive for teachers to create visual aids to explain curricular concepts, and 3) teaching Speech classes was difficult and time-intensive for teachers [25, 16]. Based on this assessment, the 2015 team also developed two software tools for the Center: *SignBook*, a "sign language dictionary creation tool with…custom video and picture capturing and categorization of entries…into topics"; and *Speak Up,* "a suite of voice-powered games aimed to familiarize pre- and partially-verbal users with the power of their voice" [16].

Setting

The Mathru Center, the setting for the most recent fieldwork, currently enrolls hearing-impaired students from 1st-6th grade. There are 57 students (39 boys and 18 girls) and 5 teachers total, with each teacher assigned a specific grade level. There is a high turnover of teachers – only one teacher stayed on between the 2015 and 2017 teams' visits. The reasons for high turn-over are related to socio-cultural factors and beyond what can be discussed comprehensively here. The first challenge facing the Mathru Center is that many of their teachers are not trained special educators, but young women; the majority of teachers who have departed since 2015 got married and left their jobs. This turnover, along with the lack of background in special education means that each new teacher learns the job while on the job – from other teachers who have the time and frequently from students. Thus, the turnover impacts both the loss of knowledge in a trained teacher and the possibility of that knowledge being passed down to the next generation

The Mathru Center had a computer lab with donated desktops, although a combination of hardware malfunctions, OS corruption, and viruses rendered most of them nonfunctioning and the school was obliged to discard them. The school also has two Lenovo laptops, donated by CMU TechBridgeWorld in 2015, that run tools developed by the 2015 team and other teaching aids the teachers have since added. Teachers use the laptops during their Computer periods, Speech periods, and/or Sign periods, each of which they have 2-3 times per week. However, due to the large quantity of curriculum material teachers must cover, periods frequently run over and teachers are unable to use the computers as often as expected. The unreliable electric supply and frequent computer viruses further prevent smooth usage of the laptops.

Speak Up! Voice Powered Game Suite

The *Speak Up! Voice Powered Game Suite* was developed, at the request of and working with teachers and administrators at the Mathru Center, because many hearing-impaired children at the Mathru Center are either unaware of the sound of their voices or have not been encouraged to explore the sounds they can make. It is an open-source[2] collection of games intended to help hearing-impaired students explore and understand their voice. As of 2015, *Speak Up* had three groups of games – Beginner Skills, Free Play, and Games – that focus on four different aspects of voice – volume modulation, pitch modulation, continuous vocalizations, and timed vocalizations. These categories were decided upon based on what the community requested, our observations of teachers' Speech classes, and the technical capabilities of the laptops' hardware and software. *Speak Up* also had a background calibration option, in which the computer asked students to be silent for 15 seconds, stored to the average volume of background noise, and ignored any sounds lower than that in subsequent games [16].

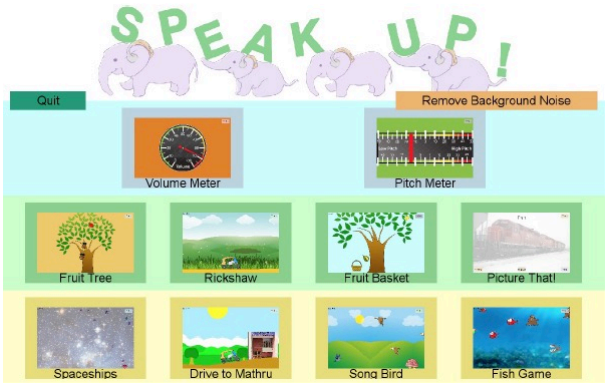

Figure 1. *Speak Up*'s main menu

[2] https://bitbucket.org/amalnanavati/istep-2015/

Beginner Skills: This group contains *Volume Meter* and *Pitch Meter*. Both games display numbers and colors that correspond to student volume and pitch, respectively. These games were intended to help students visualize and learn to manipulate different aspects of their voice [16].

Free Play: This group contains open-ended games intended to help students practice using their voice and keep them engaged in their Speech classes. In *Fruit Tree*, student vocalization makes fruits grow larger and fall off the tree. In *Rickshaw Game*, student vocalization moves an auto-rickshaw forward. In *Fruit Basket*, student volume modulation moves a basket left and right, while fruits fall from a tree above. In *Picture That*, student vocalizations gradually make completely transparent pictures more opaque. These games were intentionally made generic, so teachers could use them with a variety of curricular content and speech therapy objectives [16].

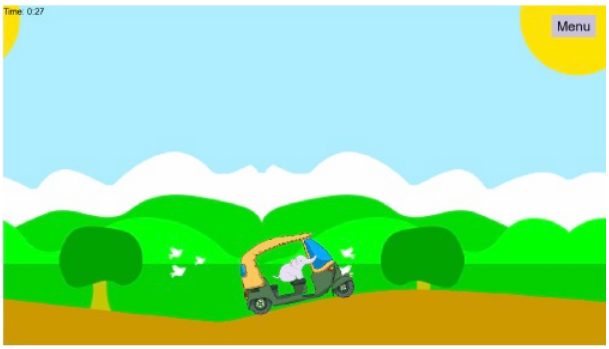

Figure 2. *Drive to Mathru*, in which students' voice propels the rickshaw forward

Games: This group contains games with explicit goals and ends. In *Spaceships*, student vocalization shoots bullets, and they have to time their sounds so the bullets hit enemies. *Drive to Mathru* is similar to *Rickshaw*, except the goal is to arrive at the students' school, and students have to continuously vocalize in order to crest hills. In *Fish Game*, student volume modulation moves the fish up and down, and the goal is to avoid obstacles for as long as possible. *Song Bird* is the same, except with pitch modulation [16].

Related Works

Over the past two decades, a number of conversational agents, games, and visualizations have been developed to support speech therapy. These tools motivate students during repetitive activities, provide information and animations on how to produce specific sounds, and/or display visuals that correspond to their vocalizations. Unfortunately, these tools are intended to be used with trained speech therapists and require complex setup and long-term usage patterns in order to be effective [7, 4, 28, 10, 22]. Tools that can be used by laypeople tend to focus on teaching students to articulate specific words and sounds. These tools analyze and provide feedback to student vocalizations by using sophisticated speech recognition algorithms and sensors; the downside is that they are expensive, slow, and/or inaccurate [12, 29, 30, 26]. Some work has been done in applying human-centered design methodologies to the development of speech therapy tools [8]. Only one of the above projects focused specifically on the speech therapy needs of hearing-impaired youth, and that project also happened to be the only one developed with and for a developing community [26]. Most research on technology for hearing-impaired individuals in developing countries focuses on tools for early assessment of the existence and degree of hearing loss [1, 17, 19]

METHODOLOGY

The most recent fieldwork that this paper details was ethnographic in nature. It included classroom observations, where I would ask a teacher for permission and subsequently sit in the back of the classroom, taking pencil-and-paper notes and/or pictures/videos on a smartphone. It also included informal conversations and semi-structured interviews with teachers and administrators: after classes; as they were walking to/from the lunch room; or in the main office, where I often sat in the principal's chair. These conversations were sometimes initiated by teachers – they asked how to use technology in particular ways or wanted to share creative teaching pedagogies they employed – and sometimes by myself. The study also included: user tests both inside and outside the classroom, where I would show teachers prototypes of additions to *Speak Up* and ask for their feedback; observations of *Speak Up* being used in classes; and trainings, both inside and outside the classroom, where I would teach one or more teachers and students on how to use the games, and stay on-hand while they used them to answer questions. I was not fluent in Kannada nor the school's local sign language; all research communications with teachers were in English, sometimes involving another teacher as a translator, and all research communications with students were translated by teachers. Teachers were aware that I was a personal guest of Ms. Muktha, founder and managing trustee of the Mathru Educational Trust for the Blind.

Since the trip was only two weeks long, the dual purposes of iterative technology development and ethnographic fieldwork were heavily interconnected. Before coming to the community, I talked with one teacher and two administrators to determine initial directions for enhancing *Speak Up*. After arriving, I spent 2-3 days observing classes and conversing with teachers, while incorporating their feedback into *Speak Up* after each school day. After creating prototypes of the modified games, I showed them to teachers and observed teachers using them in the classroom as a way of gathering feedback. I also used these observations as a way to train teachers in how to use *Speak Up* and the laptops. This continued for 1.5 weeks. The last 2-3 days were spent training teachers and administrators on how to use the games, distributing the software to the

community, working with teachers to come up with creative uses of technological tools for their classrooms, and working to ensure that the necessary systems were in place to maintain the laptops.

FINDINGS

This section details the social, cultural, and infrastructural aspects of the Mathru Center that shaped or were shaped by *Speak Up.*

Classroom Power Dynamics

The ways in which *Speak Up* was used in class both shaped and was shaped by complex power negotiations between teachers and students. This section details the ways in which authority and collective versus individual usage played a role in the school's usage of *Speak Up*.

Authority

Computers in the classroom altered the power dynamic between teachers and pupils, particularly in higher grade levels where veteran students often wrested authority from the teacher and assumed responsibility for and ownership over the games and learning. In one such class, students initially waited for the teacher to tell them what to do. However, as the class wore on, they began operating the computer at will, assuming assent unless the teacher indicated otherwise. This was partially due to age; the teachers of these classes looked young, and the students were the oldest in the school. The teachers were also newer and therefore periodically asked students for the signs of words. This upended the traditional power dynamic in which teachers are assumed to be the most knowledgeable in the class. Further, many of the students were more familiar with how computers, *SignBook*, and *Speak Up* worked than the teachers, and were vigilant while teachers used them; students would brazenly correct teachers if they made a typo or were about to click a wrong button. This precipitated teachers ceding full control of the games to students – a way of both giving them more experience with computers, and preventing further challenges to the teacher's authority.

In other classes – particular lower grade levels, classrooms where students had fewer computer skills, or classrooms where teachers had taught at Mathru for longer – teachers not only maintained authority over the class, but also recognized that authority and used it to drive particular pedagogical uses of the games. In such classes, although students might know what the teacher should do with the computer next, they would timidly provide suggestions – by pointing – without taking control of the games. This made usage of the computer a privilege that students would strive for, and they would frequently move closer to the teacher, raise their hands, make noises, or demonstrate to the teacher that they had the skills necessary to go next (especially in games where teachers required students to pronounce words), in order to next be in the spotlight. Teachers utilized this gatekeeper position by selectively deciding when students would start and end their usage of the games, and who would go next. This was influenced by students' vocal skill level, when they had last used the games, and their enthusiasm to use the games, amongst other factors.

Collective versus Individual Usage

Students' desire to be in the spotlight created a periodic ebb and flow between collective and individual usage of the games whenever teachers brought the games in front of the class. Especially with older students who were familiar with the games, they would all start using the games at once – collectively OOH-ing, for example, to move the rickshaw forward. The teacher would then silence them, and tell just one student to use the games. The class would remain silent for much of the time that student was playing the game, with the occasional student naughtily making vocalizations to test the limits of their power and independence in the classroom. However, after the student finished, other students would vie for the privilege of going next. In certain games such as *Picture That* and *Fruit Tree,* teachers were more willing to allow a collective usage of the games, since the games functionality or engagement value did not change significantly if multiple students were using it. In such cases, the teacher would focus on one student to pronounce the word correctly, but if other students tried, would work with them to a lesser degree to improve their pronunciation of the word.

Over time, teachers' focus on individuals during much of their usage of the games gave rise to a seniority amongst students, influenced by who got the most screen time, who most ardently advised the teacher on how to use the computers, who was best at the games, and who learned the games fastest, to name a few factors. These factors were in no way mutually exclusive – frequently, students with better hearing or vocal capabilities would be able to understand the *Speak Up* games better and sooner, hence be given more time with the computer, and therefore increase their computer literacy. In classes where teachers maintained more authority over the class, teachers channeled this seniority to achieve positive outcomes – they would have senior students help other students understand the games, or make the senior student a partner in computer usage by, for example, pressing a button to take a picture while the teacher held a book to the webcam. In classes where teachers maintained less authority, senior students tended to co-opt the whole interaction with the computer from other students – one such student would non-maliciously make himself the sole person who controlled the keyboard and mouse, and another student would yell loudly to show off his prowess at the sound games and drown out other students' voices. All these senior students were male.

Cultural Factors

This section details cultural factors, such as the ways in which the community found value in the games or perceived the functionality of the games, that shaped the Mathru Center's usage of *Speak Up*.

Categorization

Teachers' approach to speech therapy, and hence their usage of *Speak Up*, was premised on a binary distinction between students who "can't speak" and those who "can try [to speak]." Both groups had large variability. Students who "can't speak" ranged from students who could only make basic vowel sounds to those who could make consonant noises and mimic the timing of words, but not integrate them into full words. Students who "can try [to speak]" ranged from students who could make multi-syllable sounds that resemble simple words to those who could speak full sentences but didn't enunciate their words. Teachers similarly divided games in the *Speak Up* suite; games like *Volume Meter*, *Rickshaw*, and *Fish Game* that encourage students to "increase [their] voice" were for students who "can't speak," whereas games like *Picture That* and *Fruit Tree* that are used to encourage students to say full words are for students who "can try [to speak]."

This was in stark contrast to how we developed the games in 2015. Each game was designed to focus on one of four basic building-blocks of voice: volume modulation, pitch modulation, continuous vocalizations, and timed vocalizations [16]. It wasn't that the community did not understand those constituent components of speech. When teaching speech therapy, teachers saw it as an improvement if a student was able to time their vocalizations to align with the cadence of a word, even if the sounds they made did not resemble that word. However, teachers did not isolate that component of speech, and rather viewed it as intrinsically tied to the larger goal of saying the whole word. This related to the school's focus on teaching practical speech – speech that could be used by students to communicate outside of the school. For students who "can't speak," this consisted of teaching them appropriate uses of loud and soft voices, paired with instruction in lip reading. For students who "can try [to speak]" this consisted of teaching them to speak words that, even if not understandable in isolation, could be understood when combined with hand-gestures, facial expressions, and the context of the interaction. The goals we had encoded into the technology, therefore, was subtly misaligned with the goals the community had for speech therapy.

Due to this misalignment, teachers explicitly refocused the goals of the games to align with theirs. *Rickshaw* went from a game where hills require continuous vocalization to one where "when the road goes up, [students] must raise [their] voice." *Spaceships* went from a game that focused on timed vocalization to one where students continually vocalize in order to shoot as many bullets as possible. However, despite their re-appropriation of the games, there were indications that students were improving their ability to utilize the specific components of voice the games focused on. In 2015, one student was able to mimic her teacher's lip movements but not make any sounds. In 2017, she could not only modulate the volume of her voice, but also time her vocalizations to match the cadence of a word. However, she was unable to integrate these skills to say whole words, and teachers still put her in the category of students who "can't speak." It will take further fieldwork to understand the long-term effects on the community of technological tools that incorporate two different outlooks on how to achieve full speech capabilities: one which breaks speech down into constituent vocal skills that must be mastered, and one which focuses on speaking words in their entirety.

Mental Models

Teachers' mental models of the games, based on a combination of factors such as how they learned the games, how they taught speech, their experiences with technology, and how students wanted to use the games, created disconnects between how we had intended the games to be used and how the games were actually used. Whenever teachers were teaching the *Rickshaw* game to others, they said loud vocalizations made the rickshaw move faster. It didn't matter that the rickshaw always moved at the same speed – once students were taught that vocalizing loudly would give them the reward of speed, they would scream while using the game. Even when students realized that the games worked differently than the teacher described and began using them that way – for example, stopping vocalizations on downhills so the rickshaw could roll down by itself – they were chastised by teachers and forced to keep using the games in ways that aligned with teachers' mental models. This not only resulted in game usage that we had not expected, but also in game usage that teachers did not find particularly valuable (having students scream until their voices began to hurt); despite this, teachers did not question the games or try to find alternate uses.

Sustainability

Teacher Training

Especially for a school with high teacher turnover like the Mathru Center, institutional retention of information is crucial for the sustainability of teaching tools. Teacher trainings – which consisted of formal one-on-one sessions between new teachers and veteran teachers, periodic all-teachers meetings where teachers and administrators discussed the games amidst other topics, and informal conversations between teachers after class – served this purpose. These trainings consisted of a veteran teacher describing each game and its pedagogical purpose to one or more new teachers – frequently in terms of the binary categorization of students described above – and briefly demonstrating its usage. The veteran teacher would periodically call one or more students up – these students were the teacher's canonical examples of students in each of the two categories -- to demonstrate how they used the games and why the games are or are not useful for them. Despite this extensive survey of the *Speak Up* games, new teachers still frequently felt unprepared to use them in class, due to the disconnect between the individual use of the games demonstrated and the daunting task of using them with a larger group of students, some of whom were more familiar with the games than the teachers. Despite my

trying to set up a session for teachers to observe the authentic usage of *Speak Up* in class, it quickly turned into one of the above all-teacher meetings. It took me individually observing every teacher's class and answering any questions they had about the games in order to make them comfortable with using them in class. This warrants further research, to determine sustainable systems that researchers can use in communities with high turnover to facilitate the long-term usage of technological aids.

Figure 3. The teacher leading the all-teacher training (right) demonstrating the usage of *Volume Meter* with a student

Students as a Means of Sustainability

Students also played a large role in keeping the games sustainable at the Mathru Center. As mentioned above, they learned the games rapidly and remembered how to use them despite long (likely multi-month) gaps between uses. Even if some students forgot, the collaborative dynamic that emerged when they used the games resulted in them rapidly remembering. Further, many students knew how to operate computers, and specifically the games, better than their teachers. According to their teachers, these students had no access to technology at home; everything they learned was through using the laptops, and lessons about operating MS Paint and MS Word in the old computer lab. Because of this, and the fact that students would frequently tell teachers what to do on the computers by pointing, teachers often let students operate the games themselves. Therefore, students are crucial assets in keeping the games sustainable at the school. Although they don't have the power to dictate when the games will be used, they can ensure that information about how to use the games stays in institutional memory. Especially in the face of high teacher turnover, new teachers may be able to observe students using the games and determine what the games are for. In this way, neither students nor teachers are strictly users of the games – rather, they are partners whose crucial role in the community we must recognize and whom we must work with in order to ensure the sustainability of the tools.

Viruses

A large impediment to the sustainability of these games is viruses. Between when the 2015 team left and the 2017 team arrived, one laptop had a corrupt OS, was virus-ridden, and was missing all its networking drivers. Although our teams install anti-virus software every time we visit the community, there is no internet connection to keep the software up-to-date. We also train teachers to keep the computers virus-free and write short and simple instructions to remind them, but the school does not yet have the systems necessary to regularly maintain such technology (i.e. a dedicated person with enough time and internet access to regularly clean and update the computers). Because of how busy people at the school are, the sustainability issue is integrally tied to that of resources – time, as well as a person in the know who could rapidly fix minor issues that accrete and become insurmountable (requiring the disposal of computers mentioned earlier).

ENHANCEMENTS TO *SPEAK UP*

This section details the improvements we made to *Speak Up* during the most recent fieldwork in 2017 based on the ethnographic insights above.

Customizability

Amongst the teachers' requests for additions to the games, they repeatedly mentioned the need for games with more curriculum-specific content. For example, they suggested a game that shows a picture of a hospital, and students have to say the name of professional people found in a hospital. They suggested another game where a car is driving around town, and the student has to say the name of particular buildings in order for the car to stop there. The desire to have curriculum-specific games was inextricably linked to the concept of pronunciation; they did not just want games that related to their curriculum, they wanted games that gave feedback on a student's pronunciation of curriculum-specific terms. This was partly because, in their minds, only 2 of the 10 games were for students who "can try [to speak]", and because of the aforementioned thrust in their school to teach practical speech, which entailed focusing verbal speech therapy efforts on students with greater vocal capabilities.

Recognizing that any curriculum-specific games I develop would be temporary solutions – due to vastly different curricular needs and subject matter per grade level – I decided to create a system for teachers to add their own groups of pictures to *Picture That*. In 2015, we had incorporated a similar system into the games; teachers would take a picture (often a picture of an illustration in their textbooks) using third-party webcam software we installed on their computers, give it a descriptive name, and drag it into a folder specifically for *Picture That* images. We had even trained teachers on how to use it. Unfortunately, by 2017 teachers had not added any pictures, and knowledge of how to do so had left the institutional memory. This was partially due to the complicated multi-step process that involved familiarity with the internal file system of the computer, something that teachers did not have.

This time, I decided to use a picture-adding feature that teachers were already familiar with: *SignBook*, a sign

language documentation tool developed by our 2015 team. I modified *Speak Up* so any pictures, words, and/or folders the teachers added to *SignBook* would automatically be reflected in *Speak Up*. A big concern with any such coupling of technological tools is the dependencies it creates, and what happens when one tool fails. In order to account for this, *Speak Up* copies all of the images and metadata from *SignBook* into a local folder, so even if *SignBook* is moved, corrupted, or deleted at some later point it won't affect *Speak Up*'s usage of the words and images teachers have already added.

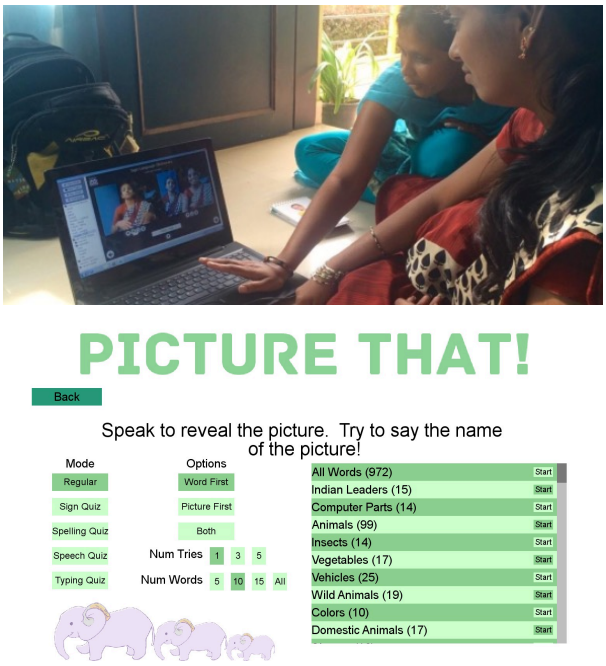

Figure 4. Top: A teacher and an administrator adding pictures and words to *SignBook* in 2015. Bottom: The new *Picture That* menu, in which teachers select from a list of folders (right) that they added to *SignBook*

Developing Technology to Utilize Teacher Expertise

The integration of *Speak Up* and *SignBook*, coupled with trainings I led to teach community members how to utilize it, addressed teachers' request for curriculum-specific content. Unfortunately, no technological solution exists for their request to have the computer grade student pronunciation. Due to the probabilistic nature of most modern speech recognition systems and the inaccuracies of phoneme recognition [2, 9, 24], having a computer understand the phonemes and words made by pre-verbal individuals (especially in a cross-cultural context where accent differs and there is large variation amongst the speech capabilities of students) is beyond the state-of-the-art. However, this gave me the opportunity to explicitly incorporate teachers' expertise into *Speak Up*, by having them press the 'r' or 'w' keys to tell the computer whether the answer was right or wrong, respectively.

This change drew teachers into the workings of *Speak Up* and allowed them to utilize their expertise to account for limitations in the technology. For example, teachers could adapt the games' difficulty to particular students, based on intimate knowledge of what constitutes improvement for that student. This was better than any difficulty setting I could have programmed into the games, which would only focus on making particular, easily quantifiable dimensions of the games more difficult – therefore simplifying the complex thought processes teachers go through when assigning difficulty levels, as well as encoding my cultural biases of what constitutes improvement in speech therapy into the games. This change in allowing teachers to control how the games grade students also allowed teachers to reflexively respond to collective versus individual uses in ways that computers are not capable of. Unlike computers, teachers could filter through multiple student voices to primarily focus on an individual. They could also dynamically change that foci when they felt other students were trying and had the potential to improve. Although this change moved particular *Speak Up* games away from stand-alone technological tools towards frameworks that require teacher expertise and cooperation in order to work, it allowed those games to achieve greater customizability, accuracy, and reflexivity than technology alone would be capable of.

Lastly, I envisioned that this change would allow the class to retain a more traditional power dynamic, where the teacher is in charge of determining when a student's performance is satisfactory enough to proceed. However, unsurprisingly this was very teacher dependent: some teachers were disinclined to grade their students wrong and merely used 'r' as a next button, whereas others saw a pedagogical purpose in grading students right or wrong and created a classroom environment in which getting an 'r' was something to strive for. This further illustrates the fact that technology is merely a tool to be used by humans; although we can design it to increase the likelihood of causing particular social effects, in the end the ways in which it interacts with a community is solely dependent on the people behind the screen.

Adapting Speak Up to the Community's Context

After observing the disconnect between teachers' beliefs about how some of the games worked versus the reality of how they worked, I changed some games to better align with teachers' mental models. For example, I changed the *Rickshaw* game such that louder voice truly speeds up the rickshaw and students must make louder vocalizations, in addition to continuous vocalizations, in order to crest a hill. I figured these changes did not take away from the goal of the *Rickshaw* game – to focus on continuous vocalization – but did make the game easier for teachers to use and train others on. However, observations of the new *Rickshaw* game revealed some students screaming as loudly as possible to move the car as fast as possible, which quickly tired out (and perhaps injured) their voice. However, getting students to speak as loudly as possible was something the teachers wanted. As opposed to changing the technology to

encourage students to moderate their voice – something that I, the researcher, felt was a better goal – I decided to leave that determination of how to use the games up to the teachers. It will be interesting in subsequent trips to see whether teachers still want students to maximize the volume of their voice, or have found alternate uses for such open-ended games.

On the other hand, with some games I explicitly pushed back against teachers' usages of the game to better aligned it with our original goal. For example, in *Spaceships* I made the game listen for spikes in volume as opposed to volume beyond a certain threshold in order to shoot a bullet. This prevented the community from using the game as they had been – continuously vocalizing in order to shoot as many bullets as possible – and required students to focus on timing their sounds appropriately. My justification for this was that there were already many games that focused on continuous vocalizations – there should be one that solely focuses on timed vocalizations. Interestingly enough, the community quickly found an additional pedagogical purpose for the new version of *Spaceships* – using it to have students practice ABCs or other brief sounds while shooting bullets.

In addition to adapting games to social factors at the Mathru Center, I also changed the games to fit their environmental context. As mentioned above, the Mathru Center is an open-air campus. This makes it near-impossible for the laptops to detect soft voices because they get drowned out by the perpetual background noise. In order to adjust for that, I enhanced the background calibration option to ask students to stay silent for 5 seconds, speak softly for 5 seconds, and then speak loudly for 5 seconds – this allowed *Speak Up* to pick up on what ranges to expect for student vocalizations, and set appropriate thresholds for the games. I also added another modality to games that relied on volume modulation, such as *Fish Game* and *Fruit Basket*, in which arrow keys control the direction of motion and volume controls the speed. This retained the focus on volume modulation, especially when students had to move the character at varying speeds. However, this also allowed students to move the character in both directions, despite background noise that previously made it hard for the character to move in the direction associated with soft voice. This change also decoupled two aspects of control of the game – direction and speed – thereby giving users more control. For example, with younger students, the teacher used the arrow keys to select a direction, and the students only provided voice. However, with older students, teachers allowed them to fully operate the game, in order to increase students' familiarity with computers.

Lastly, I adapted games to better match teachers' perceived pedagogical purposes of those games. For example, when teachers explained the usage of *Volume Meter*, they said it was useful to teach students the concept of soft voice and loud voice. However, in practice they always used it to encourage students to speak in as loud a voice as possible – this was partially due to a conflation between the numbers on the meter and a score for the game. To address this, I made a version of *Volume Meter* where stars appear along the meter and students have to steadily vocalize at that volume level in order to get that point. This explicitly changed the goal of the game from maximizing voice to modulating voice. However, because of the limitations of the Lenovo microphone, which outputted very shaky volume readings, coupled with the background noise, it was unreasonable to expect students to keep their voice at a particular volume level for more than even one timestep. Therefore, instead of changing the technology to elicit a particular effect – getting students to continually vocalize at a particular volume level – I relied on institutional memory. I spent many classes training students and teachers on how to use the game and ensured that they understood the need to keep the volume meter reading at a particular level. Although the students and some teachers understood this, it will take a future team to fully evaluate how effective the methodology of relying on institutional memory is at eliciting specific uses of the games.

Scores, Grades, and Competition

Even before arriving in 2017, teachers from the Mathru Center had been asking me to add quiz capabilities to the games. They wanted students to receive scores as a way of motivating competition and self-improvement, and they wanted the computer to tell students whether they were right or wrong. To account for this, we made a quiz mode of *Picture That*, where either the picture, word, or both appear, and students need to speak, sign, sign-spell, or type that word (these pictures and words are modifiable by teachers using the integration between *Speak Up* and *SignBook*). As mentioned above, for speaking, signing, and sign-spelling teachers have to tell the computer whether the answer is right or wrong by pressing 'r' or 'w.' For typing, the computer evaluates that itself.

These quiz modes were inspired by observations of classes (not just their Speech or Sign classes), in which teachers taught curriculum-specific vocabulary words. They would write a list of words on the board, call students up one at a time, and reflexively switch between asking them to sign, sign-spell, or say the word, depending on the particular students' skills. If the student did not get it right, the teacher would ask other students in the class. Interestingly, even though one reason teachers requested *Picture That* quiz mode was so individual students would get their scores and be motivated to improve over time, in practice the scores were never truly individual; teachers either had multiple students participate at once, or asked other students for the answer once a student got it wrong. This frequently resulted in a collaborative dynamic amongst students in order to come up with an answer – students would make suggestions or guesses, either to one another or the teacher, until they reached consensus about the answer. However, despite this collaboration, being able to use the computer to verify one's

answer was still something students strived for – especially in the typing mode, where the computer graded them. This unexpected collaborative use of *Picture That* quiz mode highlights the importance of follow-up fieldwork after introducing technological tools to a community, because community members are constantly reworking their usage of the tools due to previously undiscovered strengths and weaknesses of both the technology and the context of usage.

Based on the teachers' desire to encourage competition amongst students, I also added high score and/or best time to games that had an explicit goal. This included *Fruit Basket*, *Spaceships*, *Drive to Mathru*, *Fish Game*, and *Song Bird*. I specifically did not add a high score to games that were open-ended for fear of rigidifying the uses of the game. For example, in the original *Volume Meter*, I was concerned that keeping track of the highest number someone's voice has reached will limit the uses of the game to speaking as loudly as possible. Even though that is how the community used it, I wanted to leave other potential uses (like teaching students when it is appropriate to use which volume level) open. However, will the addition of high score shift the community's usage of the games from a collective usage to a more individual usage? It is hard to say since high score was already implemented for most of the observations in the 2017 fieldwork. However, for almost all the games in which teachers were focusing on an individual in 2017, teachers used them primarily in a collective fashion in 2015 [16]. In fact, one teacher explicitly asked me to add a time limit to the new *Fish Game* so she could give all her students an opportunity to use it, something neither us nor the teachers in 2015 even considered due to their collective usage of the games.

Robotic Car

One of the goals of our 2017 fieldwork was to investigate the potential for embodied physical agents – robots or other devices with specialized hardware and software – to enhance *Speak Up*. Prior work has revealed the potential for embodied physical agents to improve learning outcomes amongst autistic [21], deaf [23], and abled children [11]. As can be expected, the learning outcomes targeted and the consequent design of the robot is highly variable. In speech therapy, embodied physical agents could provide a directionality to speech detection, affective feedback to student vocalizations, additional ways to engage students (for example, facing the student the teacher is focusing on), or a focus on abstract curricular concepts (for example, the ability to demonstrate prepositional relationships like above and below), to name a few. However, in accordance with our focus on community-centric development, we decided to take a barebones embodied physical agent – a computer-controllable car – and work with the community to determine its potential pedagogical purposes.

To investigate this, I created a computer-controllable wireless car by dismantling the remote of a remote-controlled car and connecting an Arduino to the chip. Teachers would use a cable to connect the Arduino controller to the computer, turn on the controller, extend the antenna, and turn on the car. The car could then be moved using serial commands from the computer. In order to demonstrate to the community how the car works, I made a basic game where arrow keys dictate the direction for the car to move in, and students' voice causes the car to actually move in that direction.

The robotic car introduced the dimension of space to *Speak Up*. It separated the foci of engagement – the aspect of the games that keeps students engaged, in this case the car – from the foci of learning – the aspect of the games that focuses on teaching students and/or encouraging them to use their voice, in this case the computer. In the younger grade levels, students were simultaneously excited and frightened by the car. They would run behind it, trying to see where it would go next. They would form bridges with their legs for the car to go under, jump out of the way if the car approached them, or move the car if it got stuck under the table. However, this excitement also caused the students to make a lot of background noise – so much so that the car was effectively no longer voice activated, and solely arrow-key-operated. This suggests the potential for using the car as a classroom noise control system – students are too loud if the car starts moving.

Figure 5. A teacher (right, off-screen) controls the direction of the car using arrow keys while the student (in front of the computer) uses his voice to propel the car (orange, left). Other students (left) remain engaged by following the car

Although other teachers liked the car (there were so many shrieks of excitement from the younger classroom that they all came to see what was going on), they were concerned that it was only a toy, and wouldn't be useful for older students. Based on this feedback and further conversations with teachers, I developed a stop-and-go version of the game, in which a traffic light limits the ways in which the student can interact with the car. If the light is green, students can freely move the car using the arrow keys; if the light is red, they must answer a question correctly before being allowed to move the car again. The questions use the same image database as *Picture That* and allow teachers to decide whether to have students sign the word, sign-spell it,

speak it, or something else.

Even with the stop-and-go game, the robotic car was effective at separating the foci of engagement from the foci of learning; some students would stay close to the teacher, answering questions in the hope of getting an opportunity to operate the car, whereas others would follow the car around to prevent it from getting stuck. Although in younger grade levels the teacher alternating between operating the car herself and teaching students to operate the car, in older grade levels the students quickly became better at navigating the car than the teacher, and the teachers ceded full operational control to them. With both the basic and stop-and-go games, operating the car became a collaborative activity, with some students driving the car, others keeping track of where it was and giving the operator directions, and others creating obstacles to hinder the movement of the car. However, because of this and the fast speed of the car, all teachers were concerned about space. They all suggested that it would be better to use the car in the library – a larger, quieter room – and repeatedly asked if the car would get damaged since it was crashing against the walls of the classroom so frequently. (I responded that the car was probably tough enough to withstand it, but they should still try to limit the frequency of collisions.)

DISCUSSION

This section explores generalizable insights that came out of this work, including methodologies to navigate the power differential between researchers and community members, and recognizing and utilizing technology's capability for capacity building.

Overcoming and Utilizing the Power Differential

There was a large power differential between myself and members of the Mathru Center community: students would ask me for permission to enter the classroom I was observing, and teachers would ask me to observe their class and tell them what they were doing right or wrong. This made it near impossible to observe authentic, unobstructed uses of the games. When I asked teachers when they would be using the games in their Speech or Computer class, they said, "Right now;" not because they actually had that class, but because they were so eager to demonstrate their familiarity with the games in front of the researcher from America. As a result, most of my observations of the games being used in the classroom were orchestrated; scheduled with the teachers at the beginning of the day and occurring at a time when they should have been teaching something else. However, I still ensured that teachers were practicing the tasks they would need to do in order to use the games had I not been there: getting the laptop from the office, setting it up in the class, operating the games, etc. Further, there were still some elements of non-performative usage to the observation; the ways in which teachers were acting during the observation was shaped by and will go on to shape how they organically use the games with students.

Further, there were some unique benefits to these scheduled observations. Since no one was maintaining the façade that I was not affecting the observation, it drew down the barrier between the observer and the observed, allowing me to participate in the interaction. If teachers were confused about any aspect of the games, they would ask me immediately. If I wanted them to focus their attention on a particular new feature or game, I would ask them to. However, as soon as I had taught them what I wanted to, I stepped back and watched the teacher use it. Unlike in our 2015 trip, where observations were chiefly used for the iterative technology development process and separate teacher training occurred outside of the classroom, in the 2017 trip I brought these together, and trained teachers on games (even if they were still prototypes) while gathering observations and feedback to improve the games. This allowed me to achieve more in-depth training, give the teachers a deeper understanding of the ups, downs, and uncertainties of the technology development process, and did not hinder the observations I was able to gather to improve the technologies.

Having greater technological skills than community members also gave me more power. When we designed the games in 2015, we recognized that the power differential made it hard for us to get accurate information about how the community used them. Therefore, we built in a usage log that automatically records whenever users open *Speak Up* and play particular games. We didn't tell anyone in the community about the usage log, but left it there so we could get access to privileged information that our positioning in the community may not otherwise allow us to know. That came in handy at the beginning of the 2017 fieldwork, when all the teachers were telling me they used the games once a week each (resulting in a total of 5 times a week), but the usage logs said the games were being used about once every 2-4 weeks. The next question became: why?

I ended up telling teachers about the usage log, as a way of understanding the disconnect between their stated usage and the recorded usage. However, I was wary about this. I didn't want the teachers to feel like they might be in trouble and close themselves off to me, not did I want them to begin self-policing themselves – forcing themselves to use the games because of fear that an authority figure would know (and retaliate) otherwise. Luckily it turned out well – the teachers revealed that they had been trained in what the games do, but didn't know how to use them in a classroom setting; this prompted arranging the group teacher training, described above. Further, there is no indication that teachers remember the usage logs, that experience changed their interactions with the researchers, or that it changed their usage of the games. However, just because the use of usage logs worked well for us does not make it a methodology appropriate for all community-centric technology projects – it is rife with questions of privacy, power, policing, and the ways in which technology shapes the community. Even in this research, the true effect of using a usage log will be unclear until a follow-up research trip to the community.

Lastly, I was able to use my position of power to circulate information amongst teachers. Teachers would frequently tell me about creative methods they used to teach their students or ask me questions about how to achieve specific ends with technological tools. Although I always encouraged them to share those ideas with others, that very rarely happened, because of the pressures on teachers' time as well as the power dynamics between teachers. Therefore, whenever I had time I would share these insights with other teachers and brainstorm further ideas with them, giving credit to the original source. Although I recognize that this is not a sustainable way to encourage the sharing of creative ideas at the Mathru Center, I hope that teachers will take these examples, recognize the value of sharing innovative teaching techniques, and begin doing so themselves.

Capacity Building

When we first came to the Mathru Center in 2015, the teachers had little to no grasp of what computers were capable of doing. We would describe a game to them, and they would not be able to imagine it – we had to prototype the games in order to begin the participatory design process [16]. However, when we came in 2017, teachers were not only familiar with the games, but they were also familiar enough with the capabilities of computers to imagine and suggest other feasible games. Further, the one teacher who had been there on both trips had greatly improved her ability to operate the keyboard and mouse, and the other teachers on average were much more skilled at operating computers than the teachers in 2015. Despite not having access to computers at home, some of the older students had learned computer skills well enough to play pranks on the teacher, by using touchpad gestures to cause windows to disappear and reappear. One student had even learned basic troubleshooting skills – when *Speak Up* was malfunctioning she looked up, noticed the fan was on (and therefore creating background noise) and turned it off.

This led me to think about the ways in which I could explicitly encode our goals of boosting the community's computer literacy into the games. What if instead of viewing our games as static tools for the community to utilize, I viewed them as dynamic frameworks intended to build certain skills? These games could not jump too far ahead of their current skill level; our 2015 games allowed the teachers to modify them by moving pictures around in the file system, and they never used that feature. However, perhaps our games could stay close to the teachers' current computer literacy but keep pushing it gradually towards greater confidence when working with laptops. Or we could train older students to do some of the more complicated customization of the games, and have them serve as assistants to the teacher. In doing so, we would not just be creating static tools that have limited usefulness, but rather building up skills in teachers that can be transferred to other arenas of their work and personal lives. That was part of the motivation behind creating the ability to add pictures to *Speak Up* by using *SignBook* and having the teachers' press 'r' or 'w' to tell the games whether the students are correct or wrong.

In addition, between 2015 and 2017 teachers had developed dynamic and creative uses of the technological tools. One teacher realized that a student was having too easy of a time playing the *Rickshaw* game, so instead made him say the names of fruits and vegetables in order to move the car. One teacher had extended the usage of *SignBook* from a tool intended to document the signs of particular words to a tool used to tell stories. Another teacher used the arrow keys in the basic robotic car game to reward students for following her instructions – she would only press the arrow key when the student was making the particular sound she asked for, thereby creating the illusion that the computer recognized and approved of the student's sound. Another teacher asked me, before I had developed it, how she could add pictures to *Picture That*, and another teacher repeatedly asked me if I could teach her how to create the games so she doesn't have to wait for us to come to make her ideas a reality. This led me to explicitly start thinking about how our technology shapes self-efficacy within the community. Do we lower the community's agency if technology does everything by itself? Should we be building more open-ended frameworks, along with giving community members examples of creative ways to take advantage of those frameworks to achieve their goals?

In the time left on the trip, I focused on that. When a teacher asked me if I could provide her pictures to add to *SignBook*, I worked with her to draw the pictures instead. When another teacher asked me to create a framework for her to develop content-based quizzes, I showed her how she could utilize the *Speak Up* and *SignBook* integration to create such quizzes. When the robotic car stopped working in class, I showed the teachers how to troubleshoot it. I also spent the last day teaching a teacher Scratch, a block-based programming language [18]. Although I did not spend nearly enough time to teach the teacher to create arbitrary programs, she began seeing specific pedagogical purposes of Scratch – such as teaching the concept of bodies in motion versus bodies at rest. Once she began recognizing some uses of Scratch, I focused on teaching her how to utilize tutorials to learn Scratch herself. Although I only spent the last few days on capacity building, I hope that by having demonstrated the thought processes I go through to think of creative solutions to problems, technological or otherwise, teachers will be able to build up their potential for creativity, innovation, and self-efficacy. It will take follow-up research trip to investigate this further.

FUTURE WORK

Future work on the technical side of *Speak Up* could involve using signal processing to allow the computer to recognize student pronunciation, or using computer vision to analyze formations made by students' lips. However, such work would need to be very careful to: a) be accurate, in order to be useful and build up the community's trust in

technology; b) be flexible, in order to account for the large variation in hearing and speech skills amongst students at the Center; and c) be customizable, so teachers can easily input new content as their curriculum and speech therapy goals change over time. Other changes could include making additional curriculum-specific games, such as a version of *Fruit Tree* where student vocalization causes different professionals to emerge from a house, or real-life scenario games, such as where a student has to get someone's attention. Lastly, further games can be developed on the robotic car (such as students competing to get the car to different end goals by answering questions correctly), the robotic car's hardware can be enhanced (such as putting a sound analysis system on the car itself), and more research can be done in the pedagogical values of and cultural shifts brought about by the robotic car.

On the ethnographic front, additional research can be done in: how teachers' beliefs about technology versus the reality of how the technology works shifts over time; the long-term effects of the seniority that develops amongst students in their Speech classes; the ways in which gender intersects with technology access at the school; the ways in which students think about and perceive value in the speech games; and whether the intended social effects of our enhancements (such as incorporating a right/wrong key into *Picture That* to shift more power towards the teacher) succeeded. More work could also be done in integrating professional speech therapists' pedagogical approaches with those used at the Mathru Center – similar to the Community Advisory Boards in Neha Kumar's Projecting Health [15]. As with any community, any amount of time spent in the Mathru Center will yield interesting and invaluable ethnographic insights and questions.

On the capacity building front, the next logical step is to teach teachers and perhaps students at the Mathru Center a basic block-based computer programming language such as Scratch [18]. Scratch does not require much computer literacy to get started, but can do more if the users are more familiar with computers – this will push teachers and students to further develop their computer literacy skills. Similarly, although it is English-based, it has simple enough words that the teachers and students will understand it, and more complicated words and passages to allow them to keep learning through continued use. It also has very specific pedagogical purposes that would fit what teachers at Mathru are looking for – they can use Scratch animations to teach concepts students find difficult such as bodies in motion, respiration, prepositions, etc. Lastly, Scratch has the potential to boost teachers' self-efficacy and creativity, and such a research project would generate new knowledge on the potential for particular computer programming trainings to empower a community. We are planning to pursue this research next in our continued affiliation with Mathru.

CONCLUSION

In this paper, we described the ethnographic insights generated from two weeks of fieldwork in the Mathru Center for the Deaf and Differently-Abled and our simultaneous technology iteration on *Speak Up! Voice Powered Game Suite*. This paper highlights the need to strongly couple ethnographic fieldwork with the iterative technology development process – none of our enhancements to *Speak Up* would have been possible without an understanding of the social, cultural, and environmental contexts of the Mathru Center and how technology shapes and is shaped by them. It also stresses the need for, and utility of, studies of long-term technology deployments in developing communities – this is not only crucial for the technology development process, but also to understand and critically analyze the long term and perhaps unchangeable effects our technology interventions have in the community. Lastly, it presents results from an initial investigation into the potential for technological tools to build computer literacy, creativity, and self-efficacy within a community, discusses ways in which we incorporated that into the technology development process, and proposes directions for further research into utilizing technical tools and skills to build capacity within underserved communities.

USAGE OF FUNDS

This project was funded by a $2000 grant from IEEE Robotics and Automation Society's Special Interest Group on Humanitarian Technology (RAS-SIGHT). The proposed budget along with the actual amount spent in each category are in the table below[3].

Item	Proposed	Spent
Round trip travel between Pittsburgh and Bengaluru	$1250	$1003.51
Meals and Incidentals (ground transportation and internet)	$350	$157.38
Lodging	$300	$0
Miscellaneous	$100	$275.34
Total	$2000	$1436.23

The largest discrepancies between the proposed budget and the actual budget are in lodging and meals, since the

[3] The conversion rate from INR to USD is approximated at 64 INR to 1 USD, which it hovered around during my time in India.

Mathru Educational Trust for the Blind provided those. We were also able to find cheaper than expected airplane fares to/from Bengaluru. On the other hand, miscellaneous spending turned out to larger than expected due to unexpected technical repairs that needed to get done; reinstalling the OS and drivers on one laptop, buying a new charger for the other laptop, and buying rechargeable batteries for the robotic car. Miscellaneous funds were also spent on purchasing the Arduino, remote-controlled cars, and other parts necessary to build the robotic car. With the permission of the Chair of IEEE RAS-SIGHT, these were also used on the maintenance and repair of Stand-Alone Braille Tutors at the School for the Blind: buying and replacing fuses, buying new battery chargers, and buying new speakers. Smaller miscellaneous expenses including sponsoring a thank-you lunch at the Mathru Centre, printing, and buying CDs to store our software. Leftover funds will be used to disseminate this work at a conference.

ACKNOWLEDGMENTS

Sincerest thanks to the Mathru Educational Trust for the Blind and Ms. Gubbi Muktha, without whom this research would not have been possible. Further thanks to Dr. Priti Joshi and Xiang Zhi Tan for their help editing this paper, and Cecilia Morales for her help designing the enclosure for the robotic car controller. We also thank IEEE RAS-SIGHT and the Chair, Dr. Raj Madhavan, for generously funding this project.

REFERENCES

1. Rashid Al-Abri, Mustafa Al-Balushi, Arif Kolethekkat, Deepa Bhargava, Amna Al-Alwi, Hana Al-Bahlani, and Manal Al-Garadi. 2016. The Accuracy of IOS Device-Based uHear as a Screening Tool for Hearing Loss: A Preliminary Study From the Middle East. *Oman Medical Journal,* 31, 2: 142–145.
2. Mona A. Azim, Nagwa L. Badr, and Mohamed F. Tolba. 2016. An Enhanced Arabic Phonemes Classification Approach. In *Proceedings of the 10th International Conference on Informatics and Systems* (INFOS '16). 210-214.
3. The Braille Writing Tutor. Compassionate Engineering. 2016. Retrieved July 07, 2017 from https://compassionateengineering.wordpress.com/2016/02/26/the-braille-writing-tutor/
4. David Calder. 2008. SpeechKit: a multimedia speech tool. In Proceedings of the 10th International Conference on Information Integration and Web-based Applications & Services (iiWAS '08, 647-650.
5. Madeleine Clute, Madelyn Gioffre, Poornima Kaniarasu, Aditya Kodkany, Vivek Nair, Shree Lakshmi Rao, Aveed Sheikh, Avia Weinstein, Ermine A. Teves, Mary Beatrice Dias, and M Bernardine Dias. 2016. iSTEP 2013: Development and Assessment of Assistive Technology for a School for the Blind in India. Carnegie Mellon University Robotics Institute, 81 pages.
6. Deafness and hearing loss. 2017. World Health Organization. Retrieved July 7, 2017 from http://www.who.int/mediacentre/factsheets/fs300/en/
7. Entertaining Games and Displays for Speech Therapy. (Jan 18, 2017). Retrieved June 29, 2017 from http://www.videovoice.com/vv_fandg.htm
8. Elina Eriksson, Olle Bälter, Olov Engwall, Anne-Marie Öster, and Hedvig Sidenbladh-Kjellström. 2005. Design Recommendations for a Computer-Based Speech Training System Based on End User Interviews. In *Proceedings of the Tenth International Conference on Speech and Computers*, 483–486.
9. Santiago Fernández, Alex Graves and Jürgen Schmidhuber. 2008. Phoneme recognition in TIMIT with BLSTM-CTC. *CoRR* abs/0804.3269. 8 pages.
10. André Grossinho, Sofia Cavaco, and João Magalhães. 2014. An interactive toolset for speech therapy. In *Proceedings of the 11th Conference on Advances in Computer Entertainment Technology* (ACE '14), Article 36, 4 pages.
11. Emily Hamner, Lauren Zito, Jennifer L. Cross, Brett Slezak, Sue Mellon, Heather Harapko and Michelle Welter. 2016. Utilizing engineering to teach non-technical disciplines: Case studies of robotics within middle school English and health classes. In *Frontiers in Education Conference* (FIE 2016).
12. Sandra E. Hutchins. 1992. SAY & SEE articulation therapy software. In Proceedings of the Johns Hopkins National Search for Computing Applications to Assist Persons with Disabilities, 37-40.
13. iSTEPtbw. VLOG 1: Welcome To Mathru!. Video. (Jun 6, 2015). Retrieved July 7, 2017 from https://www.youtube.com/watch?v=ASWics-VrJk
14. iSTEPtbw. VLOG 2: Tour of the Mathru Center. Video. (Jul 5, 2015). Retrieved July 7, 2017 from https://www.youtube.com/watch?v=7PIStgUCqWY
15. Neha Kumar, Trevor Perrier, Michelle Desmond, Kiersten Israel-Ballard, Vikrant Kumar, Sudip Mahapatra, Anil Mishra, Shreya Agarwal, Rikin Gandhi, Pallavi Lal, and Richard Anderson. 2015. Projecting health: community-led video education for maternal health. In *Proceedings of the Seventh International Conference on Information and Communication Technologies and Development* (ICTD '15), Article 17, 10 pages.
16. Maya Lassiter, Amal Nanavati, Erik Pintar, Minnar Xie, Ermine A. Teves, and M Bernardine Dias. 2016. iSTEP 2015: Cross-Cultural Technology Development Towards Language Access for the Deaf and Hard of Hearing. Carnegie Mellon University Robotics Institute, 34 pages.
17. Faheema Mahomed-Asmail, De Wet Swanepoel, Robert H. Eikelboom, Hermanus C.

Myburgh, and James Hall. 2016. Clinical Validity of hearScreen™ Smartphone Hearing Screening for School Children. *Ear and Hearing*, 37, 1: 11-17

18. David J. Malan and Henry H. Leitner. 2007. Scratch for budding computer scientists. In *Proceedings of the 38th SIGCSE technical symposium on Computer science education* (SIGCSE '07), 223-227.
19. B. McPherson, M. M. S. Law, M. S. M. Wong. 2010. Hearing screening for school children: comparison of low-cost, computer-based and conventional audiometry. *Child: Care, Health and Development*, 36: 323–331. doi:10.1111/j.1365-2214.2010.01079.x
20. Meet Ms. Muktha. Compassionate Engineering. 2016. Retrieved July 07, 2017 from https://compassionateengineering.wordpress.com/2016/03/03/meet-ms-muktha/
21. François Michaud and Catherine Théberge-Turmel. 2002. Mobile Robotic Toys and Autism. In *Socially Intelligent Agents: Creating Relationships with Computers and Robots*, Kerstin Dautenhahn, Alan Bond, Lola Cañamero, and Bruce Edmonds (eds.). Springer US, Boston, MA, 125-132
22. Anne-Marie Öster. 1996. Clinical Applications of Computer-Based Speech Training for Children With Hearing Impairment. In *Proceedings of the Fourth International Conference on Spoken Language Processing* (ICSLP 96), 4 pages.
23. Becky Sue Parton, Robert Hancock, Mindy Crain-Dorough, and Jeff Oescher. 2009. Interactive Media to Support Language Acquisition for Deaf Students. In *I-Manager's Journal on School Educational Technology,* 5, 1: 17-24.
24. Phoneme Recognition. CMUSphinx. Retrieved July 7, 2017 from https://cmusphinx.github.io/wiki/phonemerecognition/
25. Shree Lakshmi Rao, Aveed Sheikh, Avia Weinstein, Aditya Kodkany, Madeleine Clute, Madelyn Gioffre, Poornima Kaniarasu, Vivek Nair, Ermine A. Teves, Mary Beatrice Dias, and M Bernardine Dias. 2016. iSTEP 2013: Exploring the Feasibility and Sustainability of Assistive Technology at the Mathru Enter for Differently-Abled. Carnegie Mellon University Robotics Institute, 27 pages.
26. Sadaf Abdul Rauf, Aatka Javed Butt, Aliza Zahid, Ayesha Jabeen, Abdul Rauf Siddiqi, and Hina Shafique. 2016. Urdu Language Learning Aid Based on Lip Syncing and Sign Language for Hearing Impaired Children. *International Journal of Computer Science and Information Security* 14, 12: 478-488.
27. Republic of India Parliament. 1995. The Person with Disabilities (Equal Opportunities, Protection of Rights and Full Participation) Act.
28. Ovidiu Andrei Schipor, Stefan Gheorghe Pentiuc, and Maria Doina Schipor. 2010. Improving Computer Based Speech Therapy Using a Fuzzy Expert System. *Computing and Informatics*, 39, 2: 303-318.
29. Chek Tien Tan, Andrew Johnston, Kirrie Ballard, Samuel Ferguson, and Dharani Perera-Schulz. 2013. sPeAK-MAN: towards popular gameplay for speech therapy. In *Proceedings of The 9th Australasian Conference on Interactive Entertainment: Matters of Life and Death* (IE '13), Article 28, 4 pages.
30. Chek Tien Tan, Andrew Johnston, Andrew Bluff, Samuel Ferguson, and Kirrie J. Ballard. 2014. Retrogaming as visual feedback for speech therapy. In *SIGGRAPH Asia 2014 Mobile Graphics and Interactive Applications* (SA '14), Article 4, 5 pages.
31. Madan M. Vaishta. 2011. Social Situations and the Education of Deaf Children in India. In *Deaf around the World: The Impact of Language*. Oxford University Press, New York, US, 352-358.
32. Madan M. Vaishta, James C. Woodward, and Kirk L. Wilson. 1978. Sign Language in India: regional variation within the deaf population. In *Indian Journal of Applied Linguistics*, 4, 2: 66-74.
33. Saurabh Varshney. 2016. Deafness in India. In *Indian Journal of Otology*, 22, 2: 73-76.