EXPLORATION OF AUGMENTED REALITY AS AN ASSISTIVE DEVICE FOR STUDENTS WITH DYSLEXIA

Abstract

Augmented Reality (AR) is a rapidly emerging technology, and its potential has not yet been fully explored. As members of Team ART, we aim to explore the use of AR as an assistive device platform for people with dyslexia, with the hopes that we could take advantage of the seamless integration of reality and computer-generated images and the attractive novelty of this up and coming platform. We began our project by surveying experts and members of the dyslexia community to determine the most helpful features and user interface for an assistive device to provide real-time feedback to users with dyslexia. Then, we developed an application on the Microsoft HoloLens to analyze users' handwritten spelling of words to provide immediate feedback. We tested the application on 19 participants in grades two through six and found that all of them improved their spelling as a result of using our device. 64.2 percent of users perceived the device to as motivating, significantly greater than the percentage of users who disliked the device. There was no significant correlation between improvement in spelling accuracy and increased motivation in regards to our device. Our novel study demonstrates that with further improvement and implementation, our application can provide assistance not only to people with dyslexia, but also to children in general.

EXPLORATION OF AUGMENTED REALITY AS AN ASSISTIVE DEVICE FOR STUDENTS WITH DYSLEXIA

Team ART

Christine Bailey, Jacqueline Deprey, Abdulfatai Fakoya, John Nolan, Ivan Quiles-Rodriguez, Cameron Soderberg, Jeong-Yoon Wu, Richard Yu

Mentor: Dr. Matthias Zwicker

Thesis submitted in partial fulfillment of the requirements of the Gemstone Honors Program, University of Maryland,

2020

Committee: Dr. Matthias Zwicker, Dr. Amitabh Varshney, Dr. Vibha Sazawal, Dr. Donald J. Bolger, Dr. Gulnoza Yakubova, Ms. Alexandra Shelton, Dr. Roger Eastman

Team ART would like to extend our greatest thanks to the many people who have helped us with our research. This includes our mentor, Dr. Matthias Zwicker, our discussants, all of the study participants, our LaunchUMD contributors, and the Gemstone Program and its staff.

Contents

1 Introduction

Dyslexia is a learning disorder which causes those affected to struggle with connecting written representations of words to the actual words themselves [6]. This difficulty leads those with the disorder to experience problems in reading and writing, even though people with dyslexia are typically as intellectually capable as their peers without dyslexia. Dyslexia is the most common learning disability, generally affecting 5 to 17 percent of children, according to various studies [22].

Dyslexia presents a major obstacle to the successful education of people with the disorder. Students with dyslexia often have difficulty keeping up with their schoolwork due to their disability [22]. Furthermore, such students often report feelings of inadequacy compared to their non-dyslexic peers, particularly arising from exclusion, ill-treatment, and negative comments from those around them [19]. It is clear that steps must be taken to ensure that students with dyslexia feel included and are bolstered with positive feedback so that they can obtain a useful and equitable education [19].

Several groups of researchers and professionals have created, tested, and implemented software vs. other methods to help dyslexic students succeed. For example, the software applications Kurzweil3000 and Sprint can help students read and write [8]. However, these applications have steep learning curves thus limiting their potential effectiveness. Other applications, such as the Prizmo application for mobile phones, have had greater success [11]. Transforming the learning experience into a game has led to mixed results [12]. New approaches, including those based on the emerging technology of augmented reality (AR), could provide different or greater benefits to students with dyslexia.

There are several reasons to believe that an AR-based approach could alleviate the problems posed by dyslexia. First of all, AR offers the possibility of real-time, hands-free feedback. Compared to other technologies, AR can allow students to perform their work with minimal technological interruptions. Furthermore, several studies have shown that using AR in an educational context helps to motivate (i.e. stimulate interest and energize) children [25, 27]. Increasing children's excitement about learning can cause them to put more effort into and be more persistent about their education, so an AR-based approach could help students succeed even more than a conventional, less exciting approach could [16] .

Motivated by these concerns, we set out to answer the following research question:

> What features of an AR application are most effective for improving the spelling abilities and increase the motivation of people with dyslexia?

To answer this question, we surveyed members of the dyslexia community (dyslexia researchers, special-aids teachers, people with dyslexia) about what features they believed would work well in an AR application created to help students with dyslexia succeed in the classroom. Using the results of this survey, we created a prototype AR Technology application (called ART) to correct handwritten spelling mistakes. We tested ART on children to determine whether or not ART helped children improve their spelling ability and whether or not ART increased these students' motivation to learn.

Users of the application can write text on a piece of paper and then use the device to take a photograph of their writing, process and correct the text, and then present the corrected text on a head-mounted display.

ART was created based on survey results from members of the dyslexia community We designed ART to enable users to obtain real-time feedback on spelling and handwriting errors. We hope that users will develop firm reading

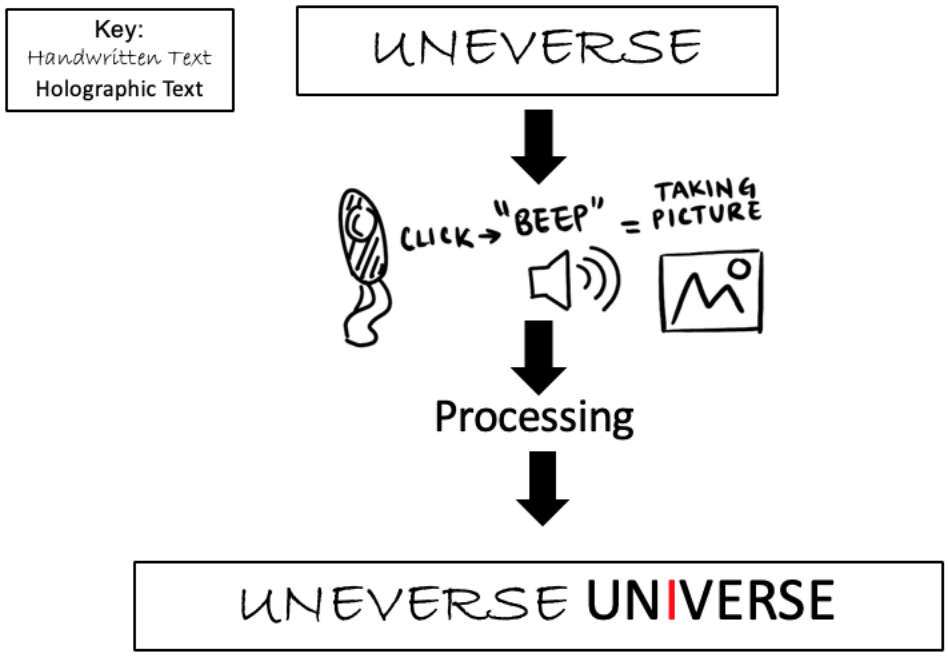

Figure 1: Diagram of the steps taken by the ART application. The top box contains the user's view through the Hololens which is a misspelling of the word "universe". Then the user would click the Hololens clicker and then there is a beep that communicates the device has taken a photo. After that the photo is processed and the correction for the word is displayed in the user's line of sight

and writing skills as a consequence of using ART.

To determine the effectiveness of ART, we invited children to participate in one-hour testing sessions using ART. During these sessions, children were tasked with spelling (in handwriting) a series of age-appropriate words. After spelling each word, the children were instructed to use ART to correct their spelling. After re-testing, we found that ART both improves students' accuracy in spelling and motivates them to learn more.

As we built our application, we ran into a few challenges to keep in mind when developing future iterations of assistive programs such as ours. One challenge we faced in developing ART was the unreliability of our initial spell-checking. In particular, the amount of noise (irrelevant background information) introduced by text recognition was often enough to cause the spell-checking program to misinterpret the written text. To circumvent this, we restricted ART to use only the words which we were using in the tests. We ensured that these words were all different enough that the chance of ART confusing these words was negligible. This restriction was deemed worthwhile for a proof of concept. However, a larger-scale application would need to be able to correct text from a significantly larger subset of the English language.

Another challenge that arose in developing ART was the incompatibility of many pieces of software crucial to the successful operation of ART. Because AR is still an emerging technology, many of the tools needed to implement it are still in development. When these tools are modified and updated, they may lose or gain the ability to work with other software ART relies on. This makes it difficult to create a consistently working version of ART. After several iterations, we found versions of these software that work together consistently; we describe these in Section 3 of this paper.

2 Literature Review

2.1 Dyslexia

About one in five people have a language-based learning disability, the most common of which is dyslexia, which is characterized by difficulty in decoding, reading, spelling, and recognizing words [7, 30]. One major common criterion for dyslexia diagnosis is that the reading accuracy of the patient is greater than 1.5 standard deviations below the mean [15]. However, once the difficulty in reading has been established, one must differentiate between two potential problems: dyslexia (which encompasses the problem of learning to decode print), general problems in reading comprehension, or even another learning disability. Problems of decoding are related to difficulty in oral language development and lack of speech-sound phonological skills.

In order to be a successful reader, one must be able to integrate all the regions of the reading circuit in the brain that command language, visual information, and orthographic processes with accuracy and speed [20]. Previous research has discovered dyslexia-susceptible genes that contribute to the making and efficiency of the reading circuit in the brain [20]. There are proposed regions of the brain that make up the reading circuit: the temporal-parietal cortex, the superior temporal cortex, the occipito-temporal cortex, and the inferior frontal cortex [22]. Each have its own role in terms of integrating orthographic and phonological information, specializations for print and rapid word processing, speech planning, semantics, comprehension, and other general cognitive functions needed for reading. Although the timing of the specialization of function of the regions are disputed, their roles in reading are clearly solidified. Researchers have found lack of activation in these regions in dyslexic patients [22]. When genes that contribute to the proper functioning of these brain regions are knocked out in mice, there are consequences in the mouse that pertain to dyslexia [20]. For example, removing the DCD2 gene leads to impairment in visuospatial memory, visual discrimination, auditory processing, referencing, etc. [20].

As there is substantial evidence that dyslexia is tied with brain and gene development, early diagnosis of dyslexia may become possible as research progresses, thus allowing for even earlier intervention. However, it is currently not feasible for all children to be diagnosed with dyslexia until they are at the age where the symptoms become obvious. This can deprive many children of early intervention promoting decoding and training in phoneme awareness needed for those with dyslexia, which can cause even greater issues for children with dyslexia later in life. The lack of reading instruction specifically focused on helping students with decoding and phonemic awareness can also attribute to disparities in a person with dyslexia's reading and comprehension abilities. If such instruction was instituted for all children, it would help all children with these skills, but it would also be particularly helpful for children with dyslexia.

Learning disabilities make education more difficult and therefore present an impediment to the success of those with such disorders. Any method to alleviate the symptoms of dyslexia would aid many people. Since most treatments for dyslexia use educational tools to help enhance reading [9], the use of AR as an aid could benefit the affected population.

2.1.1 Technological Approaches

Creating effective technology to assist people with dyslexia is difficult, especially because the definition of dyslexia is vague and many people who have dyslexia are also diagnosed with one or more other learning disabilities [29]. The technologies and strategies that individuals use to assist themselves vary from person to person.

One case study describes the different technologies that university students with dyslexia have used to cope with their disability [24]. Each student surveyed had issues with spelling, reading, and writing. All of the students used speech-to-text software to help with their writing, and two of the students used mind mapping software to keep their thoughts structured and ordered before writing.

Another study had participants all use the same existing software – Kurzweil3000 and Sprint [8]. These programs were reading and writing-based and gave audio or visual feedback when requested by the students. While these assistive technologies have potential benefits, the majority of participants found that the time it took to learn the new technology as well as to scan their materials in order to use the software outweighed any benefits they had.

The textbook *Access all Areas: Disability, Technology, and Learning* [23] lists the main situations where people with dyslexia typically fall behind their peers: listening, reading, organization / memory, written language, and calculations. The textbook recommends recording lectures or giving students with dyslexia preprinted overviews of lectures, using colored overlays, highlighting sections of text, and using mind mapping and / or speech to text software. The textbook also notes that many spellchecking applications do not work as desired because the spelling mistakes people with dyslexia make are typically due to phonetic errors, and the checker has difficulty correcting such misspellings.

Multisensory computer-based training was used in another study with children who had developmental dyslexia and children who did not [18]. The children were given computer training with the goal of recoding text into audio and visual codes. All children who participated in the computer writing training improved their writing skills, and even their abilities to write words that were not included in the training improved as well.

Another study used a device featuring automatic speech recognition to help facilitate immediate reading intervention [17]. The study found that for people with dyslexia, especially with reading, immediate feedback in combination with a multi-sensory interface is key. This study informed our design decision to make our application provide immediate feedback.

Other studies demonstrated that learning devices lead to an increase in motivation for students with dyslexia. One such study used an app equipped with optical character recognition and speech synthesis as 'reading glasses' for students with dyslexic profiles [11]. Students were asked to read pseudowords (fake words with examples of different phonics) over around 20 sessions per child. The researchers demonstrated positive results – the students' skills increased more quickly than expected. Students, regardless of their performance, were more motivated and interested when using the app, as it was so novel compared to their usual means of study. Even if a device might not quantitatively make a student better at reading or writing, any increase in motivation can drive that student to put in a greater effort.

Another study looked at the use of gamification to motivate students with dyslexia [12]. The researchers used classDojo, a popular platform for teachers, to award badges to students for their achievements and to update parents and other classmates of the badges earned. The main findings of the study were that each student reacted differently to the amount of badges learned, especially as not all teachers enlisted in the study gave the same number of badges. They also found that many students with dyslexia are naturally less motivated because they do not see themselves making as much progress as their peers. While these students all shared a common learning disability, they did not all react similarly to the reinforcement provided by the badges. It is important to acknowledge that there might not be one technology or strategy that will work for all people with dyslexia, and while there are some techniques that help a broader population, there is no “best method.”

2.2 Augmented Reality

Several current developments in technology motivated Team ART’s decision to use AR as a platform for our application. The large amount of money being funneled into technology-based learning makes technology-based solutions to learning disabilities more viable than ever. AR devices, and specifically AR headsets, allow users to receive immediate feedback on their writing by allowing them to view computer-generated corrections projected onto the natural world, as depicted in Figure 2. In addition, AR headsets allow users to receive feedback without needing to hold a device in their hands, so that users can transition seamlessly between writing and obtaining corrections.

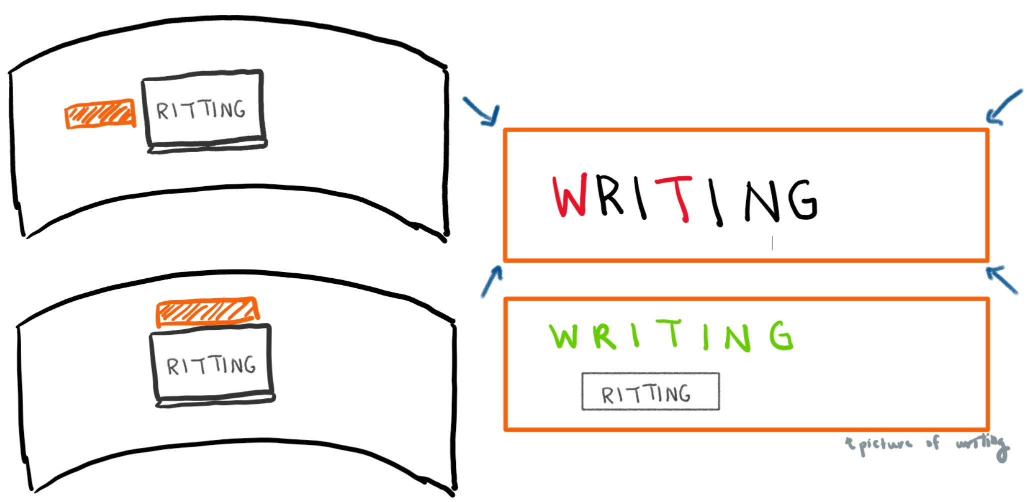

Figure 2: A visualization of how AR can superimpose text and other features onto real environments.

AR offers numerous benefits to potential learning applications. According to [33], AR has the potential to create authentic learning environments, help students learn subjects that otherwise require real world experience, and to motivate students. Furthermore, [31] states that AR has the ability to provide learners with a sense of immersion and give them a sense of ubiquitous learning. AR can also help users to learn content in 3D perspectives and visualize the invisible [31].

AR has already been experimented with as a tool to help children in the classroom. One study created an AR protractor to help elementary students with geometry, and found that all of the students that tested their device thought their device was helpful to them, and increased their motivation to learn geometry [25]. Another study created several prototypes of AR applications focused at aiding elementary students in various mathematics curriculum areas [26]. After testing on recruited teachers, the study concluded that AR applications could be well suited to helping elementary students with a wide range of mathematical topics [26].

Yet another study created an AR book designed to help preschoolers with reading, and found that it helped students identify letters [27]. This study also found that the combination of real world and virtual objects grabbed the students' attention, generating excitement, engagement, and enjoyment [27]. Overall, the current implementations and studies of AR technologies in learning environments have indicated AR to be a useful and motivational tool in the classroom. These studies have also shown the wide range of applications that can utilize AR, and they help pave the way for even more educational AR applications.

2.3 Current AR Options

At the time our team selected an AR platform, there were two main, widely available AR headsets on the market: the Microsoft HoloLens and the Meta 2. We describe the devices in this subsection, deferring the reasons for our ultimate selection of the HoloLens to Section 3. An image of a child wearing a HoloLens is provided in Figure 3 to illustrate the physical specifications of current AR devices.

Figure 3: Photograph of a child using the Microsoft HoloLens. Used with permission.

The Meta 2 features 90 degree field of view, hand and positional tracking, microphones, and 720p RGB front facing camera for video [21]. It weighs about 1.1 pounds and requires a cable to be attached to a computer in order to function [21]. The Meta 2 is light but also relatively bulky in size, due to the large amount of room required above the eyes to hold the hologram projectors [21]. Even given the adjustable head strap and exchangeable forehead pads, the headset is unbalanced, and can become uncomfortable over long periods of time [21]. The holograms look realistic, but are not particularly bright [21]. While the device occasionally fails to recognize gestures correctly, the holograms are relatively easy to manipulate [21].

In contrast, the Microsoft HoloLens does not require a tether to a computer, which aids users' mobility when using it. The HoloLens features several high definition cameras, as well as depth sensors for environment mapping and video capture, microphones, and a hologram field of view of roughly 30 degrees [14]. The display has a maximum of 720p resolution for holograms, going down to a minimum of 320p, at a minimum of 60Hz refresh rate as well as minimum and maximum distances for hologram placement [13]. It supports gaze, gesture and voice tracking for controls and weighs 1.27 pounds [13]. The weight of the HoloLens can cause high pressure on the nose piece, which can lead to discomfort or pain. The HoloLens also supports remote operation over Wi-Fi, allowing for recording of holograms or remote manipulation of the holographic environment [13]. Although the holograms on the HoloLens are brighter and appear more solid or lifelike than those of the Meta 2 (due to more holographic density), the field of view on the HoloLens is small, at 30 degrees, which can make the user to occasionally miss seeing holograms [13].

2.4 Handwriting Recognition Technology

A successful handwriting correction application must be able to recognize text written by the user in order to identify errors and provide feedback. We researched several different libraries and approaches in our attempts to find solutions to this fundamental problem. The options we considered fit into three broad categories: manual, open-source, and commercial. We discuss these options here.

We considered designing our own neural network for text recognition. Neural networks are large, multilayered mathematical constructs which can be used to classify or label new data when "trained" on a sufficiently robust data set. They have been used to solve many real-world problems, including the label-ing of photographs, analysis of text, and classification of biomedical relations, among countless others [28, 32, 34]. One major advantage of neural networks is their versatility. Given enough time and training data, a neural network could be set up to categorize or recognize data from practically any domain. However, this usually requires large amounts of time and training data needed to achieve acceptable results.

As an open source option, we considered a combination of the libraries OpenCV and Tesseract. OpenCV[1] is an image processing library that provides computer vision and machine learning functionality to real time processes. OpenCV offers a library of algorithms which implement computer perception and machine learning capabilities such as facial recognition, producing point clouds, image stitching and more. OpenCV also features image processing such as blurring or contrast adjustment. We considered using the optical character recognition (OCR) engine Tesseract[2] in conjunction with OpenCV. Tesseract is a minimal open source library that provides OCR functionality. In a test case, we found that Tesseract recognized printed text accurately but required the input images to be processed (to remove noise) using OpenCV or a similar application.

The last family of options we considered was commercial cloud-based text recognition services, e.g. Amazon Web Services[3], Google Cloud[4], and Microsoft Azure[5]. These options require minimal training or preprocessing on the user's end, but, in exchange for this simplicity, the user sacrifices some degree of control over the recognition process. Furthermore, as opposed to the other options considered, these services typically charge a fee for services.

[1] Available at `https://opencv.org`
[2] Available at `https://github.com/tesseract-ocr/tesseract`
[3] Available at `https://aws.amazon.com`
[4] Available at `https://cloud.google.com`
[5] Available at `https://azure.microsoft.com`

3 Methodology

The goal of our project was to create an AR application that can provide fast and effective spelling corrections to students while they are writing. In pursuit of that goal, we sought information from researchers and experts in fields related to education and learning disabilities, as well as from the Dyslexia community itself. The information sought from experts was gathered through in-person interviews with various individuals on-campus. With this information, along with suggestions from experts, we were able to develop a survey targeted at individuals who either have dyslexia or have experience interacting with individuals with dyslexia. The results of this survey would then be used as an important component when determining the components to use when building our application.

3.1 Survey

Before beginning the development of our application, we began surveying the Dyslexia community to gain a better understanding of what setbacks they experience when learning and how those could be prevented through the use of AR. We define the Dyslexia community as individuals who have dyslexia themselves or interact with persons who have dyslexia (teachers, doctors, parents, etc.) on a regular basis. We chose this population since we believe they would have first or second hand experience with the struggles we aim to identify regarding how dyslexia impacts a person's ability to read and write. By understanding current methods for dealing with dyslexia, our team aimed to hone in on what the community liked about current methods and what could be improved. The questions asked on the survey and the results can be seen in the following sections.

3.1.1 Initial Survey

1. Do you have Dyslexia?

2. Do you work with people who have dyslexia on a daily basis?

   (a) If yes, what is your relationship with them? Select all that apply:

      - Parent
      - Teacher
      - Doctor
      - Other (please specify)

3. (If 1 answered affirmatively)

   (a) Do you have struggles reading?

      - (If yes) What helps you read better? Select all that apply:
        - Changing word font
        - Changing text size
        - Changing color of paper
        - Changing text color
        - Other (please specify)
        - No changes can improve my reading performance

   (b) Do you have struggles writing?

      - (If yes) What helps you write better? Select all that apply:
        - Live handwriting correction
        - Predictive text suggestion
        - Auto correct
        - Grammar correction

    - Other (please specify)
    - No changes would improve my writing performance
  - (If any changes selected)
    - Where would you like to see the corrections?
      * Above the written word
      * On top of written word
      * Below written word
    - When would you like to see corrections?
      * While writing
      * After writing

(c) Do you think a wearable assistive device to help with reading / writing would be helpful?

(d) Would you be comfortable trying a wearable assistive device to help with your reading / writing struggles?

4. (If 2 answered affirmatively)

(a) Do the people / person with dyslexia you work with have struggles reading

  - (If yes) What would help them read better? Select all that apply:
    - Changing word font
    - Changing text size
    - Changing color of paper
    - Changing text color
    - Other (please specify)
    - No changes can improve their reading performance

  - Unsure

- Do the people / person you work with have struggles writing?
  - (If yes) What would help them write better? Select all that apply:
    * Live handwriting correction
    * Predictive text suggestion
    * Auto correct
    * Grammar correction
    * Other (please specify)
    * No changes would improve their writing performance
    * Unsure
  - (If any changes selected)
    * Where would you like to see the corrections?
      · Above the written word
      · On top of written word
      · Below written word
    * When do you believe they would like to see corrections?
      · While writing
      · After writing
- Do you think a wearable assistive device to help with reading / writing would be helpful to the people / person you work with?
- Do you think the people / person you work with would be comfortable trying a wearable assistive device to help with their reading / writing struggles?

3.1.2 Survey Results

Out of the 28 survey respondents, four people stated that they had dyslexia while the rest had experience working with people with dyslexia. The response from the individuals with dyslexia was not conclusive towards any specific assistive device design, due potentially because of a lack of participants. The respondents who worked with people with dyslexia had a wide range of responses, but in general reported that hearing the word out loud helped with reading the most, and feedback or text correction would be the most helpful for people with dyslexia when writing.

Notable responses and expert opinions:

- “Feedback needs to be more direct. No language like ’did you mean’. Just have the correct text.”
- “Give feedback at end of sentence and every minute.”
- Important to establish proof of concept, something that demonstrates the feasibility of the device: “Only need to prove increase of student motivation/engagement and potential increase in ability for it to be implemented.”
- Introducing the device to the children: “You should put on device first to show them how it’s done. Hand device to parent to let them do it / increase trust in the kid. Hand device to kid and let them do it for themselves, instead of trying to fit them. This process will increase their acceptability of the device as a whole.”

Answer	Percentage	Count
Yes	14.29%	4
No	85.71%	24

Table 1: Question: Do you have Dyslexia?

Answer	Percentage	Count
Yes	100.00%	4
No	0.00%	0

Table 2: Question: Have you struggled with reading either in the past or present?

Answer	Percentage	Count
Changing word font (please specify what changes)	16.67%	1
Changing text size	16.67%	1
Changing color of paper	0.00%	0
Changing text color	0.00%	0
Hearing the text out loud	66.67%	4
Other (please specify)	0.00%	0
None of these changes have improved my reading performance	0.00%	0

Table 3: Question: What helps you read better, or has helped you read better in the past? Select all that apply.

Answer	Percentage	Count
Yes	50.00%	2
No	25.00%	1
Unsure	25.00%	1

Table 4: Question: Has isolating words or sentences on a page ever improved your reading ability?

Answer	Percentage	Count
Yes	100.00%	4
No	0.00%	0

Table 5: Question: Do you have struggles with any aspect of writing?

Answer	Percentage	Count
Vocabulary	9.09%	2
Grammar	13.64%	3
Spelling	18.18%	4
Swapping characters	18.18%	4
Writing characters incorrectly (i.e. writing letters backwards or upside down)	18.18%	4
Mixing up similar sounding words	18.18%	4
Other (please specify)	4.55%	1

Table 6: Question: Which of the following aspects of writing do you struggle with / have struggled with in the past? Select all that apply.

Answer	Percentage	Count
Live handwriting correction	11.11%	1
Predictive text suggestions	11.11%	1
Spelling auto-correct	44.44%	4
Grammar corrections	33.33%	3
Other (please specify)	0.00%	0
None of these would improve my writing performance	0.00%	0

Table 7: Question: Do you think any of the following would help you write better or would have been useful to you in the past? Select all that apply.

Answer	Percentage	Count
Yes	50.00%	2
No	50.00%	2

Table 8: Question: Do you think a wearable assistive device to help with reading / writing would be helpful?

Answer	Percentage	Count
Yes	75.00%	3
No	25.00%	1

Table 9: Question: Would you be comfortable trying a wearable assistive device to help with your reading / writing?

Answer	Percentage	Count
Yes	75.00%	21
No	25.00%	7

Table 10: Question: Do you interact with people who have dyslexia on a regular basis?

Answer	Percentage	Count
Yes	95.00%	19
No	5.00%	1

Table 11: Question: Do the people / person with dyslexia you work with have struggles reading?

Answer	Percentage	Count
Changing word font	14.29%	5
Changing text size	17.14%	6
Changing color of paper	5.71%	2
Changing text color	5.71%	2
Hearing text out loud	40.00%	14
Other (please specify)	14.29%	5
None of these changes will improve their reading performance	0.00%	0
Unsure	2.86%	1

Table 12: Question: What would help them read better? Select all that apply.

Answer	Percentage	Count
Yes	64.71%	11
No	5.88%	1
Unsure	29.41%	5

Table 13: Question: Would isolating words or sentences in a page help improve the reading ability of individuals with dyslexia that you work with?

Answer	Percentage	Count
Yes	100.00%	18
No	0.00%	0

Table 14: Question: Do the people / person that you work with have struggles writing?

Answer	Percentage	Count
Vocabulary	15.94%	11
Grammar	17.39%	12
Spelling	26.09%	18
Swapping characters	13.04%	9
Writing characters incorrectly (i.e. writing letters backwards or upside down)	10.14%	7
Mixing up similar sounding words	17.39%	12

Table 15: Question: Which of the following aspects of writing do the individuals with dyslexia you work with struggle with the most? Select all that apply.

Answer	Percentage	Count
Live handwriting correction	10.71%	6
Predictive text suggestion	23.21%	13
Auto correct	17.86%	10
Grammar correction	17.86%	10
Audio feedback	19.64%	11
Other (please specify)	7.14%	4
No changes would improve their writing performance	0.00%	0
Unsure	3.57%	2

Table 16: Question: What would help them write better? Select all that apply.

Answer	Percentage	Count
Above the written word	100.00%	1
On top of the written word	0.00%	0
Below the written word	0.00%	0
None of these, would prefer audio feedback	0.00%	0

Table 17: Question: Where do you believe they would like to see the corrections?

Answer	Percentage	Count
While writing	55.56%	10
After writing	44.44%	8

Table 18: Question: When do you believe they would like to see / hear the correction?

Answer	Percentage	Count
Yes	93.75%	15
No	6.25%	1

Table 19: Question: Do you think a wearable assistive device to help with reading/writing would be helpful to the people/person you work with?

Answer	Percentage	Count
Yes	75.00%	12
No	25.00%	4

Table 20: Question: Do you think the people / person you work with would be comfortable trying a wearable assistive device to help with their reading / writing struggles?

3.2 System Design and Implementation

Based on the feedback received in the survey, we decided to create an assistive device to help individuals improve their writing by providing visual and audio feedback along with real-time corrections. Although all respondents with dyslexia also reported having trouble with reading, we decided it was more feasible to tackle the one issue. Since we decided to focus more on helping the individuals with their writing, we looked more thoroughly at the responses to questions about writing to determine criteria for our application. Before the analysis, our basic goals were broken up into; finding the best platform to build the application on, the problems they experienced when writing, and the most effective way to fix those problem. During our analysis, we found that in terms of application platform, 50% of the "dyslexia" respondents said a wearable assistive device would help them, and 94% of "other" respondents said it would help as well. When asked about what issues the respondents had with writing, the majority of responses indicated spelling and mixing up words that sounded alike. Finally, when asked for what would help with their writing, responses were split between automatic text correction, predictive text suggestion, audio feedback, and grammar correction. Once we had all these criteria set, our task was to find a proper platform to develop the application on, begin creating the application itself, integrate text recognition and correction as well as the criteria above into the application, and finally test it.

Figure 4: An image of the HoloLens with motherboard separated. This includes the central processing unit which executes all computer instructions, a graphics processing unit that renders the images seen on the HoloLens, and finally a Holographic Processing Unit which processes logic for displaying the actual holograms and super positions[14].

Figure 5: An image of the IMU of the HoloLens which includes an accelerometer, gyroscope, and magnetometer, four environment cameras, an RGB camera, four microphones, and a Time Of Flight camera. This allows the device to collect a wide array of information to accurately track user movements and information about its surroundings. The HoloLens uses the IMU to implement the mapping of the physical would as well as other features such as minor hand tracking for gesture recognition, head tracking, and even sound localization to simulate sound coming from different positions on the mapped space[14].

Figure 6: An Image of Holographic lenses that generate graphics. This includes the holographic display which is responsible for projecting the graphics of the text, menus, transitions, and overall user interface on the device. The display reflects the digital light that the device sends to it and combines it with the natural light it is receiving to make an accurate image of the computer generated images being shown to the user[14].

3.2.1 Application Platform

One of the most important aspects of the criteria was how we would provide feed back, as this was what would actually be helping the individuals using the application. This was why we decided to use AR as the basis for the application, because it provides additional solutions for feedback that other traditional technologies did not. In its simplest definition, AR is technology that allows for the visual superimposition of computer generated information in the users vision. This is especially useful in our case because we wanted to provide feedback to individuals as they write, and would like it to be done almost instantly. The feedback would appear to be a physical object in the users physical environment, because it is in their vision, and could even be potentially interacted with to engage the user even more. This was the different approach we planned to take in hopes of creating something that could help individuals where traditional means could not.

AR capable devices and hardware are mainly available in two sources: handheld devices and headsets. The choice of which platform to develop on was important since it sets the stage for how users were going to interact with our application. For instance, due to their ubiquity, mobile devices was one of our top choices for developing the application. In addition to this, both iOS and Android devices had toolkits that allows for the easy building of AR experiences onto mobile platforms. However, we wanted our device being something that could provide immediate feedback, and would not hinder the individual writing in anyway. Having to hold a phone would cause an obvious conflict with that, and was not further considered. As the "hands-free" experience was more critical in designing the project, the other available option were AR headsets. At the time, our feasible choices for headsets were between the Meta 2 and the Microsoft HoloLens.

In the end we chose the Microsoft HoloLens as our AR platform. This is mainly due to hardware differences between the two, with one major difference being that the HoloLens is a standalone device. Rather than having it be hooked up to a computer to siphon power and operate, the HoloLens has a built-in battery which impedes the user even less. In addition to this, the HoloLens hardware contains three main components; the motherboard, inertial measurement unit (IMU), and holographic display which can be seen in Figures 4, 5, and 6. When working together, these components allow for quick and accurate measuring of the physical world space not found in other devices. This is the crux of AR in general, and directly impacts how well the user experiences the application. In addition, the HoloLens is equipped with various sensor, cameras, and an advanced display system, which allows for better quality of feedback we want to provide. Figures 4, 5, and 6 have images and further explanations for different parts of the Hololens hardware.

The HoloLens was also more popular than the Meta 2, so many people had already implemented applications for it. This meant that we had a clear reference starting point, and the device was well documented with some amount of developer support. The device was also backed by Microsoft so there was very little chance that support for the product would be dropped during the time we were working on the project. Overall, we decided to use the HoloLens because it was more practical, powerful, and had more documentation and resources available. With the platform selected, we then moved on to finding software to actually make the application.

3.2.2 Unity Environment

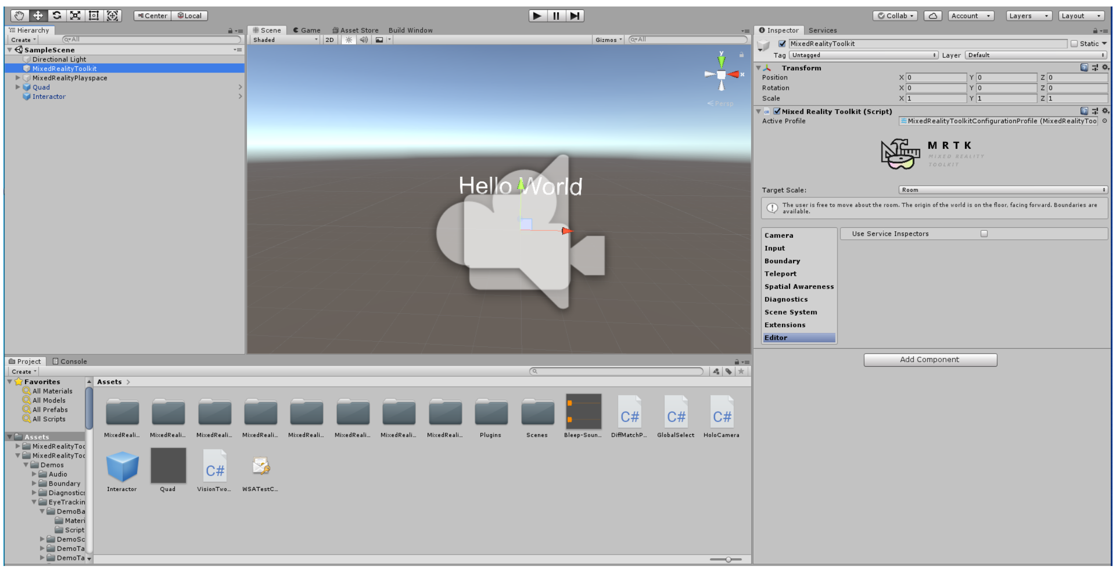

Figure 7: Image of Unity environment used to build the application

We decided to use a video game engine as the basis for our application because of the robust features included in such engines. These engines come with built-in tools to create AR experiences, including resources for managing object physics, graphics, user interfaces, and more. While an AR headset like the HoloLens already contains AR capable hardware for processes such as gesture recognition, environment mapping, among with other crucial procedures, a game engine enables developers to utilize these features to suit their needs and integrate even more features. As such, game engines are generally the

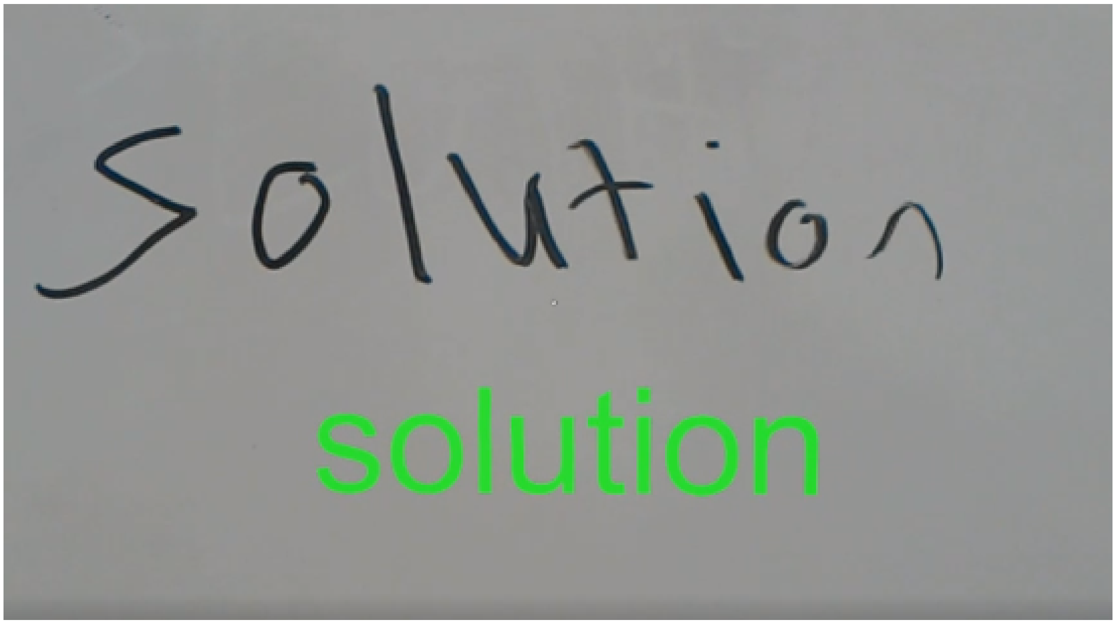

Figure 8: Image of ART application when user writes “solution” correctly. The word is displayed in green to show that it is correct.

most widespread way to develop AR applications.

The HoloLens development process is most compatible with the Unity engine, which motivated our choice of Unity as opposed to other engines. Most sample AR applications we researched used Unity as their development engine, so it also stood out because of its popularity in commercial use. Furthermore, as we were conducting research, we found that most tools and documentation for HoloLens development assumed developers were using Unity.

Another motivating factor of our use of Unity was the Mixed Reality Toolkit (MRTK). The MRTK, released by Microsoft, allows developers to take advantage of common features and libraries when developing for different AR or VR headsets. Standard features in the MRTK include mapping the environment, placing virtual objects, recognizing gestures, and localizing sounds. The MRTK is also easy to use – when building to the HoloLens, developers only need to drag and drop the MRTK asset into the Unity software, and choose the HoloLens as the specific headset used. With that MRTK included into the project, all basic features for a HoloLens AR application would be added. Once this was done, the Unity environment was set and ready for the building of an application.

3.2.3 Text Recognition

When deciding on an approach to implement text recognition, our initial idea was to train our own neural network to perform OCR. To train such a neural network, we would need a large amount of images of text coupled with what they contain as a string of text. These examples would be split into two different groups: one for training and one for testing.

While training, the neural network would take an image from the training set and attempt to produce the corresponding string of text. If the text produced by the network was incorrect, the neural network would update its parameters via a process known as backpropagation. This training would go on for some time, after which the network would be tested on the data in the testing set to see how it performs on novel images and to prevent the algorithm from overtraining on the training data set. This process is complex, and several fields of research have been working on these problems for decades.

We decided not to pursue creating our own neural network for a variety of reasons. First of all, given the importance of text recognition as a general task in computer vision, we anticipated that publicly or commercially available software would be effective enough for the task at hand. Such software would typically be improved over time by its developers, so our application would also improve over time. Secondly, we expected that writing our own neural network would introduce undesirable dependencies or requirements into the application. This could lead to complications preventing the application from functioning. It takes time and money to train an effective neural network, and we did not expect that our current resources and knowledge would be sufficient enough to create a network that serves our purposes better than pre-made options.

Because creating our own neural network was not a viable option, we considered two popular open source OCR frameworks; OpenCV and Tesseract. The open source computer vision library OpenCV is able to not only detect handwriting, but also perform important pre-processing on images. For example, OpenCV could perform edge detection which is used to detect a piece of paper within an image. This functionality would make it easier to determine the text within the bounds of a piece of paper. Tesseract was also open source, but was designed specifically for OCR. So given a photo containing text, both software would perform pre-processing to reduce noise in the image and potentially extract the text.

However, we decided that neither of these tools would be effective for our application. Tesseract was only available in the programming languages C++ and C, and OpenCV was only available in C++ and in a Unity-incompatible port to C#. We expected this to present many build and integration difficulties with the C# and Unity-based workflow favored by the HoloLens. Furthermore, we anticipated that the HoloLens would not be able to pre-process images and perform OCR effectively due to low computational power. It was possible that using a remote server to perform pre-processing and OCR would allow us to bypass this difficulty. However, this option would add another level of complexity to the application, making development and execution of the application more difficult. As a result, we chose not to pursue this option.

Once we realized that the above approaches would not work for our application, we decided to look at commercial software. We expected that these services, being published by well-established technology companies, would offer high-quality suites for text recognition and related problems. The particular software we sought to use was Microsoft Azure, together with its Cognitive Vision Services. We tested the text recognition functionality of this service, and found that it produced results which were clear and accurate enough to be used for our application.

The main draw towards Azure was that its Cognitive Vision Services was available as an application programming interface (API) - an abstraction that allows two programs to communicate with one another. Because of this, we are able to perform image processing and data extraction by sending requests to the API, eliminating the need to perform processing on the HoloLens or on our own server. From there the API connects to Azure's own server, which performs the requested tasks and sends a response with the needed data back to the HoloLens. This process avoids the potential concerns over the computational power of the HoloLens and completes quickly, which would still allow us to provide timely feedback. Furthermore, Azure's libraries are available in C#, making it easy to integrate into the Unity application. All of these factors helped us decide to use Azure's Cognitive Vision Services in our application.

3.2.4 Text Correction

Similar considerations to those surrounding text recognition arose when we investigated methods for performing spell-checking and text correction. We had a similar list of options for this task: writing our own software, commercial software, and then open source software. Creating our own software for spell-checking, much like creating a neural network for OCR, seemed as if it would not be worth the effort. Hence we only seriously considered using commercial or open source software.

When we looked into commercial software available for text correction, we saw that Azure contained a library designed for spell-checking. Since it was another API developed by Microsoft, it had all the same benefits as the Cognitive Vision Library, and seemed as if it would work well with the application. However, when developing the application, we found that this was not the case.

Azure's spell-checking API was optimized for correcting phrases rather than individual words. When given a sentence, it would analyze the context to see where mistakes were made and properly correct them. However, when given a single word without context, Azure might fail to properly correct an incorrectly spelled word. Because we wanted to test students' spelling of individual words rather than sentences, we needed to find another option.

After some consideration, we decided to restrict the words our application would recognize to a hard-coded list of words to be used in testing. These words were chosen after researching challenge words for each grade, and were selected from sources tested by the International Reading Association [1, 2, 3, 4, 5]. We decided to choose challenging words to encourage the users to make mistakes so that we could test the efficacy of our application. These words were selected to limit the similarity of the words to each other in order to improve the accuracy of the recognition. This prevented the application from incorrectly recognizing words, and made the application work consistently from student to student. Although this approach does not generalize well, we considered it valid as our goal was not to create an application for widespread distribution, but rather to present a proof of concept.

In order to perform this text correction, we used the DiffMatchPatch[6] tool. DiffMatchPatch is a an open source tool containing several useful func-

[6] Available at `https://github.com/google/diff-match-patch`

tions for plain-text manipulation, and in our case, specified text identification. The "Match" portion of the program takes a word and a length of text, then attempts to find the closest match to the word in that text. We found that the C# version of DiffMatchPatch was compatible with our application, straightforward to use, and effective at finding matches. Consequently, we incorporated DiffMatchPatch into our application.

3.2.5 Application Design and Composition

With all decisions on the different components of the application finalized, the process of building the application was split into two parts. The first part being the design of the user interface, and the second part being getting the core logic of the application to work with the libraries and commercial software selected. When designing the interface, the most important aspects were to make it simple as to not distract the user, and to make the experience as comfortable as possible while providing the necessary feedback. In order to mitigate distractions, we made the application with two sources of feedback: holographic text and a beep noise for audio. The purpose of the holographic text was to display the original word if correct, or corrections if the word was wrong, with the beep acting as an indicator for various stages of the application and also as the important source of audio feedback as recommended. We also made it possible to drag the the text around the world space so it could be easier for the user to view the text.

While the user interface was being developed, we also created the code for correction and recognition. These were developed independently and on separate code bases for the overall goal of increasing productivity. As stated previously, in order to implement the recognition we utilized the Azure API. Using that API along with further examples provided by Microsoft, a script was written that would capture text. On a high level, the script used a photo image that was taken on the HoloLens as input. From there, the API authorized us as a legitimate user, and the aforementioned image was sent as a request to be processed by the Azure server. Once finished, the server would then send a response back to us with a variety of information extracted from the image. We then parsed through that information to extract the recognized text.

After obtaining a string containing the recognized text, we passed the string to DiffMatchPatch for processing. The specific implementation uses five word lists that were separated based on grade level. Each of which contains all of the potential words that will be matched. In order to select the appropriate grade level before the matching process, we set up five unique key words corresponding to grades 2-6; alpaca, baboon, crocodile, dragon, and elephant. We chose relatively obscure words as keywords to prevent the accidental triggering of a grade level in the midst of testing. These keywords were then integrated with the HoloLens' microphone capability. Since there was only a specific subset of words that need to be matched, it was reliable enough to simply match the words extracted from the image to the word list. This allowed for a high rate of accuracy in returning the correct word from the image taken under the conditions of our study.

Once these parts were finished, we had three separate components; the Unity Scene that the user of the application would experience, a script that could recognize text given an image, and a second script that when given text would respond with corrections. The composition of all these parts was done in Unity due to its ease of adding functionality. All we needed to do was to drag the pre-written scripts into Unity, and have it be referenced by a main script. The first functionality of this script was to initialize the camera on the

HoloLens to be ready to take an image. Once the image was taken, it was passed on as a parameter to the text detection code that was referenced. The detection returned a string output which was sent to the correction script that was being referenced. Finally, the result was referenced by the text object in the Unity scene that the user would see, which was then updated to display the text correction returned. Thus, the application was complete.

3.3 User Study

After building out our device, we wanted to test the question of whether or not our application was successful in helping people with dyslexia. Our study was defined around two key metrics of success: the first being whether or not the students testing our device perceived the real-time feedback as helpful or not, and the second being if they felt more motivated to learn as a result of using our device. This was based on the survey, which asked questions about motivation and helpfulness of the device. These metrics are derived from our goal to increase the retention of these students in schools as previously defined in our literature review.

To answer this question, we decided to test on children in grades two through six since our research indicated that this is primarily when students with dyslexia begin to struggle with spelling. Our previous research indicated that tools designed to help students with dyslexia read and write better also help students without dyslexia. As a result, we opened this portion of our testing to any student, including those who did not have dyslexia, who fell within that grade range to increase the number of test participants we could recruit during the limited time frame of this project.

For our testing procedure, participants began by being guided by a team member in basic usage of the Microsoft HoloLens; how to fit the device on the head, how to operate the application, and basic usage guidelines. To assess the efficacy of the writing assistance provided by the application, a short phrase was read out loud by a member of our team to the participant, who then was asked to write said phrase on a whiteboard. They were then asked to use their best handwriting and all capital letters. Test participants were also told to take their time writing since it was important that the handwriting was neat for our application to be able to interpret what they wrote.

When a user would load our application they were greeted with some welcome text. The user would then say one of the key words to select the grade level out loud, which was detected by the HoloLens' microphone in order to select the proper word list for the grade level. The user would then point the Hololens at their text, and take a photo using either the Hololens' built in "airtap" gesture with their fingers or by using the physical clicker device that was paired with the Hololens. When the image was taken, the user would hear a beep, alerting the user that the image was taken. The user would then hear a second beep when the corrections were ready and the display would change from the welcome text into the feedback from the program. If the word was completely correct, the word would be displayed as green, and if not, the subset of the word that needed correction would be red, with the rest of the text in white. After they received the feedback, they were asked to erase the whiteboard and say "next" when they were ready for the next word. A member of our team recorded what the participant wrote on the board. Participants were given a set of 10 words for this procedure, selected from suitable challenging word lists.

After the test participants ran through this procedure with each of the ten words, they took a break before continuing on to repeat the same process with the words they spelled incorrectly in the first round. By retesting the test participants on the same words, we were able to see if the real-time spelling corrections helped the students with their writing in the short-term.

After the students finished the second round of writing words on the whiteboard and receiving corrections, they were given a series of post survey questions which they were to respond to verbally if they agreed or disagreed with the statements. This post-testing questionnaire was used to learn more about the users' experience throughout testing. Our research showed that students with dyslexia not only have a harder time in school with reading and writing, but they also have tend to fall behind because of lower motivation levels to learn.

Our aim of this questionnaire was to determine if the students had a positive or negative experience with this device and if it motivated them to want to continue to write and learn. We assessed motivation by asking questions relating to whether the students perceived the device as helpful, liked using the device, and/or wanted to use the device more. To remain unbiased, we asked the test participants an equal number of positive and negative statements regarding different aspects of their experience with the device which we used to gauge their overall feelings towards it. We used binary choices for answering these questions to force participants to take a stance on this device. Furthermore, we wanted to make sure the younger children we tested on could understand the questions and answer appropriately based on how they were feeling, so we limited the number of choices they had to each question. To assess the increase in motivation, we would assign each positive statement a participant selected a score of $+1$ and each negative statement a score of -1. We then planned to calculate a score for each test participant and analyze the results.

As long as we can prove that our AR application motivated students to learn, we can conclude that it is helpful in the academic setting, even if there is inconclusive evidence that it improves the students' spelling abilities.

4 Results

4.1 Data

After finishing our testing, we began tallying the percent increase of our test participants to see how much they improved as a result of using our device. All of our test participants improved from the first iteration of writing the words to the second. The amount of words our test participants spelled correctly (out of a total of 10) in the first and second rounds can be seen in Table 4.1, where each participant is represented by a row in the table. A summary of the change in number of words spelled correctly can be seen in Table 4.1.

Round 1	**Round 2**
8	10
7	10
5	8
8	10
8	10
7	10
1	3
2	5
0	1
8	10
1	4
6	8
6	10
5	7
2	4
5	6
3	5
7	9

Table 21: Number of Words Spelled Correctly in Round 1 vs. Round 2

In addition to collecting data on how the students' performed using our device, we were also able to gather quantitative data from our survey given to students after testing the device. We decided to use the term "device" to be

Number of Words Improved	Number of Participants
1	2
2	10
3	4
4	2

Table 22: Changes in Number of Words Spelled Correctly from Round 1 to Round 2

all inclusive of the hardware of the headset itself as well as our application. Because the students are interacting with the technology as an all inclusive experience and their developmental neurological state, experts in the field of educational feedback suggested that we refer to the two components together as a collective item.

To draw meaningful conclusions from our data, we focused on statements that related to positive and negative statements about the test participants' interactions with the device.The format of the data is as follows:

Survey Sentence – Number of Children that agreed

Positive statements included

1. "This device makes learning more fun" – 12
2. "This device makes me want to try writing more" – 9
3. "I want to try more things with this device" – 14
4. "I like wearing this device" – 10
5. "I think this device is helpful" – 16

Negative statements included

1. "I do not like using this device" – 1
2. "This device is fun, but not for schoolwork" – 4

3. “It is hard to wear this device” – 5

4. “I do not think this device can help me” – 2

5. “This device is distracting” – 1

By breaking up our survey questions based on positive and negative feelings towards the device, we were able to quantify a test participant’s overall experience using our device. Having a numerical score for each child allowed our team to perform a variety of statistical tests on the results and calculate the impact our device had on increasing the motivation of students to learn.

4.2 Analysis

We first began by analyzing whether our application was statistically significant in improving the spelling performances of our test participants. To do this, we preformed a paired-tailed T-test on the data with $H_0 : \mu = 0$. The true mean difference in students’ scores from the first round to the second is equal to 0, so the device has no effect on a student’s short-term spelling abilities. This null hypothesis was compared against our alternative $H_\alpha : \mu > 0$. The true mean difference in students’ scores from the first round to the second is greater than 0, so the device has positive effect on a student’s short-term spelling abilities. This resulted in $t = 2.777$ with 17 degrees of freedom and $p = 0.006455$.

Because our p value is less than our significance level of 0.05, we have enough evidence to reject our null hypothesis in favor of our alternative that an augmented reality assistive device that provides spelling correct to students does increase the short-term spelling abilities of students.

This conclusion is not surprising though, for anyone given immediate feedback and asked to repeat the same task a short-while later will most likely recall at least some of the details from their first iteration. What our team was more interested in, was how our augmented reality assistive device affected the motivation of students.

Before we began performing our statistical analysis of the motivational component of our data, we took a broad overlook to see what aspects of the device worked and what were areas that could be improved.

Upon looking at the preliminary data, we can see that our device is perceived as helpful since that statement had the greatest number of responses at 16. Our device is also seen as engaging for “I want to try more things with this device” and “This device makes learning more fun” had the second and third greatest number of respondents with 14 and 12 respectively. Because the positive statements had an average selection rate of 64.2%, which was much greater than the average selection rate of 13.7% for negative statements, it seems like our device was perceived as helpful.

To begin our statistical analysis, we first wanted to assess the overall feelings towards our device, as it pertained to using an augmented reality headset for learning as a whole. To do this, we assigned each positive statement a participant selected a score of +1 and each negative statement a score of −1. We then calculated a score for each test participant which resulted in the data set of Table 4.2.

+5	+3	−4	+1	+4
+3	+1	+4	+4	+2
+4	+4	+5	+2	+4
+4	+5	+4	+3	

Table 23: List of scores indicating participants’ feelings about the device

We then preformed a one-tailed T-test on the data with $H_0 : \mu = 0$. The true mean score for students using our device equal to 0, so the device has no effect on a student’s perception of learning. This null hypothesis was compared against our alternative $H_\alpha : \mu > 0$. The true mean score of students using our device is greater than 0, so the device as a whole does motivate students to learn. This resulted in $t = 6.354$ with 18 degrees of freedom and $p < 0.0001$. Because our p value is less than our significance level of 0.05, we have enough evidence to reject our null hypothesis in favor of our alternative that an augmented reality assistive device that provides spelling correct to students is perceived as a helpful learning tool.

In addition to checking if our device was helpful as a whole, our team was also interested in the relationship between a child's improvement between the two trials as well as their feeling score toward the device. To assess this, we performed a χ^2 test on this data which resulted in the following breakdown:

Improvement	Motivation		
	-4 to 2	3 to 5	Total
1 to 2	4/18	8/18	12/18
3 to 4	0/18	6/18	6/18
Total	4/18	14/18	18/18

Table 24: χ^2 Table

Because all of the students who tested our device improved from the first trial to the next (with the exception of one test participant who was unable to do a second trial due to device malfunction), we decided to focus this relationship on how much they improved. Improvement was measured by the increase in the number of words the student spelled correct from the first to the second trial, while motivation was measured by the student's survey score.

This resulted in $\chi^2 = \sum i = 14\frac{(O_i - E_i)^2}{E_i} = \frac{1}{7}$ and $p = 0.706$ with one degree of freedom. Because our p-value is not smaller than our significance level of 0.05, we do not have sufficient evidence to reject our null hypothesis that there is a relationship between a student's improvement level and their feelings towards the device.

Through testing our application on 19 test participants and running several statistical tests on our results, we are able to make three key conclusions. First, an AR application giving live feedback statistically increases the performance of students' spelling. Second, an AR application can significantly increase the motivation in children to learn. Third, the increase in motivation of a student using augmented reality is not linked to how much one improved in their spelling.

5 Future Work

This project's goals were to test an assistive, augmented reality, application for children and gauge their engagement level and responsiveness when learning. We expected for students who used the application would see improvements in their spelling. However, we did not expect to be able to obtain comprehensive results on spelling improvements due to time constraints which limited our ability to perform a longitudinal study. Therefore, we chose to collect feedback based on the engagement level of the children when writing with the device, and children's willingness to continue to use the device for learning purposes. This decision was motivated in part by the fact that it is much faster to test motivation than long-term improvement in spelling. In addition, because of cost constraints, we could not distribute devices to families or teachers, so performing testing in a more "natural" environment was not possible.

The results of our study pointed to increased motivation in children who used the application. One of the primary purposes of our study was to determine if students' were more motivated to learn when they used assistive devices. According to the Department of Education national technology plan, technology such as the HoloLens can accelerate, amplify, and expand effective teaching practices [10]. Our study corroborated this claim by suggesting that the HoloLens can be an effective motivational agent for children with learning disabilities such as dyslexia. More broadly, this study indicated that AR provides a novel approach that could help children engage better with certain subjects. It's possible that AR applications like the one described in our study could be used in schools to help children with disabilities improve their skills.

In further work on this project, we would like to pursue two major aims. Our first goal would be to improve the generality of the application. In particular, the application could be improved by enabling it to interpret words that do not come from a fixed list. Our application used for testing could only interpret words from a hard-coded list. Though we created a version of the application which tried to interpret and correct words from the totality of the English language, we found it returned too many errors to conduct a meaningful study on its efficacy as a learning tool. Improvements in optical character recognition, improvements in spell-checking technology, and improvements in computer vision could all help generalize our application, which would be paramount for its application in the field.

Additional functionality within the application could also be useful for generalized use. For example, adding positive reinforcement stimuli (e.g. an animation when a word is spelled successfully) could improve the application's motivational capabilities. Many other beneficial features could be added on to the basic application presented here. Or, instead of just correcting the errors, the application could track the errors to create a body of data on what errors students make to help teachers identify patterns, focus lesson plans for an individual, and track improvement.

Finally, decreases in latency of the application are also needed to make it useful as a generalized learning tool. Possible improvements towards this end include implementing an algorithm to identify what new text has been added since the last time a user asked for feedback. Currently, the application performs OCR and spelling corrections on all the text in an image every time a user asks for feedback, which increases the delay between feedback and feedback request. Limiting the processing to only new text thus would decrease the delay. Additionally, running OCR and spell-checking operations on the device itself rather than outsourcing these tasks to a cloud-based service would decrease the delay. Finally, some users may want to have real-time feedback without having to request it every time, so an improved application could give the user an option to receive constant real-time corrections rather than having to click a button to request feedback.

Our second goal would be to run longer studies on the effect of AR technology on students (especially students with dyslexia or other learning disabilities). Such long-term studies could have control groups and regular surveys to test whether or not these technologies have persistent effects on student educational outcomes. Both motivation and spelling ability improvements could be measured in these long-term studies. This could help schools determine whether or not to invest in these technologies. Such long-term testing would also help to confirm or deny the results of this study. As it stands, it is unclear if our device improved motivation simply because it was novel to students, and if the effects of this novelty would decrease over time.

6 Conclusion

Through this project, we created the ART application to help students with dyslexia. We evaluated the efficacy of the application and determined that it improves users' excitement about learning. However, as per the results of our study, we cannot yet determine whether the relationship between this increase in motivation and improvement during the test is significant. However, our results as a whole suggest that AR can help students with dyslexia succeed in the classroom.

While completing this project, we learned valuable lessons about the current AR landscape and applications of AR to education. We found that AR is still in its infancy, which can make developing applications difficult. For example, because the features and compatibility of AR software libraries are still in flux, it can be difficult to combine various libraries to create a working application. Furthermore, many current AR devices (including the HoloLens) are not designed with children in mind and can be large and heavy enough that children have difficulty using them.

Despite these difficulties, further work on applications of AR with respect to dyslexia could be useful. It would be interesting to see whether or not a software application that allows users to apply visual filters to their writing (e.g. changing the perceived color of the paper) helps users improve their reading or writing. Another interesting avenue of research would be extending the functionality of the application (or creating a similar application) that is able to correct a significantly larger subset of the English language and can therefore be applied to real writing rather than the spelling-test scenario presented in this paper. Finally, testing the application or a similar application on a larger body of students over a period of months or years would help verify whether or not AR is actually useful in changing educational outcomes.

7 Appendices

7.1 Post Testing Questions

Think about how you felt testing our device and then please select the statements that you agree with:

- I do not like using this device
- This device makes learning more fun
- This device makes me want to try writing more
- This device is fun, but not for schoolwork
- It is hard to wear this device
- I want to try more things with this device
- I do not think this device can help me
- This device makes me want to try reading more
- I like wearing this device
- This device is distracting
- I think this device is helpful
- I like reading and writing on my own
- This device can help with school

Results

1. "I do not like using this device" – 1
2. "This device makes learning more fun" – 12

3. “This device makes me want to try writing more” – 9
4. “This device is fun, but not for schoolwork” – 4
5. “It is hard to wear this device” – 5
6. “I want to try more things with this device” – 14
7. “I do not think this device can help me” – 2
8. “This device makes me want to try reading more” – 8
9. “I like wearing this device” – 10
10. “This device is distracting” – 1
11. “I think this device is helpful” – 16
12. “I like reading and writing on my own” – 11
13. “This device can help with school” – 11

7.2 Testing Procedures

- Recruitment : We recruited children in grades 2-6 to participate in our study. We did this distributing our flyer advertisement to the Parent Teacher Student Associations of local elementary and middle schools in Maryland. We also posted information about our study on social media and by emailing teachers who specialize in these grades. Our recruitment efforts will be targeted at Prince George’s County schools for they are the closest in proximity to the University of Maryland where our testing will take place.

7.3 Building the Program

1. Clone the repository (`https://github.com/Team-ART-Gemstone/ArtProject`)
2. Open the project Using Unity 2018.4 LTS. Potentially newer versions will also work.
3. Install MRTK version 2.0 from the Microsoft Github page `https://github.com/microsoft/MixedRealityToolkit-Unity`
4. Player settings should have mixed reality supported.
5. Update Azure ComputerVision API key with your own in Assets/VisionTwoHelper.cs
6. To build select Universal Windows Platform, Target Device Hololense, x86 Architecture, D3D Build Type, 10.0.10240.0 Minimum Platform Version, as well as have Unity C# projects checked.
7. Build to App folder.
8. Open Visual Studio Sln.
9. Install Nuget Packages
    - Windows UWP
    - Azure.CognitiveServices.Vision
10. Build to your Hololens. Find the IPv4 address and enter it when prompted after selecting remote device under build. To install a debug version select Release or Debug build, otherwise select Master Build to install a regular version.

Glossary

API Application programming interface. A specification of functions, classes, or procedures through which one computer program can make use of another. 37, 38

AR Augmented reality. A collection of technologies allowing users to superimpose virtual images and computerized features on their real-world surroundings as viewed from a headset or other electronic device.. 6, 7, 9, 11, 14–16, 19, 31–35, 43, 50, 51, 53, 54

MRTK Mixed Reality Toolkit. A collection of software tools used in developing for the Microsoft HoloLens.. 34, 35

OCR Optical character recognition. A field of computer science concerned with automatically detecting and transcribing text.. 18, 35–37

References

[1] *2nd Grade Vocabulary Word List*. URL: https://www.flocabulary.com/2nd-grade-vocabulary-word-list/ (visited on 2020).

[2] *3rd Grade Vocabulary Word List*. URL: https://www.flocabulary.com/3rd-grade-vocabulary-word-list/ (visited on 2020).

[3] *4th Grade Vocabulary Word List*. URL: https://www.flocabulary.com/4th-grade-vocabulary-word-list/ (visited on 2020).

[4] *5th Grade Vocabulary Word List*. URL: https://www.flocabulary.com/5th-grade-vocabulary-word-list/ (visited on 2020).

[5] *6th Grade Vocabulary Word List*. URL: https://www.flocabulary.com/6th-grade-vocabulary-word-list/ (visited on 2020).

[6] American Psychiatric Association et al. *Diagnostic and statistical manual of mental disorders (DSM-5®)*. American Psychiatric Pub, 2013.

[7] *Attention-Deficit/Hyperactivity Disorder (AD/HD) and Dyslexia – International Dyslexia Association*. URL: https://dyslexiaida.org/attention-deficithyperactivity-disorder-adhd-and-dyslexia/ (visited on 10/03/2017).

[8] Nadia Diraä et al. *The Use of ICT to Support Students with Dyslexia*. Springer, Berlin, Heidelberg, 2009. ISBN: 978-3-642-10307-0.

[9] *Dyslexia - Treatment Overview*. (Visited on 2017).

[10] Office of Educational Technology. *Reimagining the Role of Technology in Education*. 2017.

[11] Linda Falth and Idor Svensson. "An app as 'reading glasses' – a study of the interaction between individual and assistive technology for students with a dyslexic profile". In: *International Journal of Teachnig and Education* 3.1 (2015), pp. 1–12.

[12] Daniel Gooch et al. "Using Gamification to Motivate Students with Dyslexia". In: ACM SIGCHI. San Jose, May 2016, pp. 7–12.

[13] Mark Hachman. *We found 7 critical HoloLens details that Microsoft hid inside its developer docs*. URL: `https://www.pcworld.com/article/3039822/we-found-7-critical-hololens-details-that-microsoft-hid-inside-its-developer-docs.html` (visited on 03/02/2016).

[14] *HoloLens (1st gen) hardware*. Sept. 15, 2019. (Visited on 10/28/2019).

[15] Charles Hulme and Margaret J. Snowling. "Reading disorders and dyslexia". In: *Current Opinion in Pediatrics* (2016), pp. 731–735.

[16] Melissa Hurst. *The Importance of Motivation in an Educational Environment*. 2012.

[17] Husniza Husni and Zulikha Jamaludin. "ASR technology for children with dyslexia: Enabling immediate intervention to support reading in Bahasa Melayu". In: *US-China Education Review* 6.6 (June 2009), pp. 64–70.

[18] Monika Kast et al. "Computer-based multisensory learning in children with developmental dyslexia". In: *Restorative Neurology and Neuroscience* 25.3-4 (2007), pp. 355–369.

[19] Monicca Leseyane et al. "Dyslexic learners' experiences with their peers and teachers in special and mainstream primary schools in North-West Province". In: *African Journal of Disability (Online)* 7 (2018), pp. 1–7.

[20] S Mascheretti et al. “Neurogenetics of developmental dyslexia: from genes to behavior through brain neuroimaging and cognitive and sensorial mechanisms”. In: *Translational psychiatry* 7.1 (2017), e987.

[21] Lucas Matney. “Hands-on with the $949 mind-bending Meta 2 augmented reality headset”. In: *Tech Crunch* (Mar. 2, 2016). URL: `https://techcrunch.com/2016/03/02/hands-on-with-the-949-mind-bending-meta-2-augmented-reality-headset/` (visited on 10/28/2019).

[22] Ola Ozernov-Palchik and Nadine Gaab. “Tackling the ‘dyslexia paradox’: reading brain and behavior for early markers of developmental dyslexia”. In: *Wiley Interdisciplinary Reviews: Cognitive Science* 7.2 (2016), pp. 156–176.

[23] Lawrie Phipps, Allan Sutherland, and Jane Seale. *Access All Areas: disability, technology and learning.* JISC TechDis Service and ALT, Jan. 2002. ISBN: 1-904190-18-9.

[24] Geraldine Price. “Creative solutions to making the technology work: three case studies of dyslexic writers in higher education”. In: *ALT-J, Research in Learning Technology* 14.1 (Mar. 2006), pp. 21–38.

[25] James Purnama, Daniel Andrew, and Maulahikmah Galinium. “Geometry Learning Tool for Elementary School using Augmented Reality”. In: IEEE. Bali, Indonesia, Apr. 2014, pp. 145–148.

[26] Iulian Radu, Betsy McCarthy, and Yvonne Kao. “Discovering educational augmented reality math applications by prototyping with elementary-school teachers”. In: IEEE. Greenville, SC, Mar. 2016.

[27] Dayang Rohaya Awang Rambli, Wannisa Matcha, and Suziah Sulaiman. “Fun Learning with AR Alphabet Book for Preschool Children”. In: Procedia Computer Science. 2013, pp. 211–219.

[28] Waseem Rawat and Zenghui Wang. "Deep convolutional neural networks for image classification: A comprehensive review". In: *Neural computation* 29.9 (2017), pp. 2352–2449.

[29] Gerd Schulte-Korne. "The Prevention, Diagnosis, and Treatment of Dyslexia". In: *Dtsch Arztebl Int* 107.41 (Aug. 2010), pp. 718–727. DOI: 10.3238/arztebl.2010.0718.

[30] *Statistics*. URL: http://www.dyslexiacenterofutah.org/dyslexia/statistics/ (visited on 10/03/2017).

[31] Hsin-Kai Wu et al. "Current status, opportunities and challenges of augmented reality in education". In: *Computers & Education* 62 (2013), pp. 41–49.

[32] Zichao Yang et al. "Hierarchical attention networks for document classification". In: *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*. 2016, pp. 1480–1489.

[33] Steve Chi-Yin Yuen, Gallayanee Yaoyuneyong, and Erik Johnson. "Augmented reality: An overview and five directions for AR in education". In: *Journal of Educational Technology Development and Exchange* 4.1 (2011), pp. 119–140.

[34] Yijia Zhang et al. "Neural network-based approaches for biomedical relation classification: A review". In: *Journal of biomedical informatics* (2019), p. 103294.