The Future of Engineering Education
2024 Annual Conference & Exposition
Oregon Convention Center
Portland, OR . June 23 - 26, 2024
ASEE
Paper ID #44324

Tactile Learning: Making a Computer Vision Course Accessible through Touched-Based Interfaces

Dr. Seth Polsley, University of Nebraska, Lincoln

Dr. Seth Polsley is an Assistant Professor of Practice in the Jeffrey S. Raikes School of Computer Science and Management at the University of Nebraska–Lincoln, with his academic home in the School of Computing. His research focuses on the combination of intelligent systems design and human-computer interaction in order to support novel educational and universal computing experiences.

Ms. Amanda Kate Lacy

Amanda Lacy is a PhD student at Texas A&M University in the department of Computer Science and Engineering. Her interests are broad, with an emphasis on applying computing to promote access to information and spaces, both virtual and physical. She hol

Samantha Ray, Texas A&M University

Samantha Ray is a Computer Engineering PhD student at Texas A&M University. Her research focuses on creating intelligent systems for tasks that require human-like levels of understanding. She has previously worked on human activity recognition (HAR) syste

Dr. Tracy Anne Hammond, Texas A&M University

Dr. Hammond is Director of the Texas A&M University Institute for Engineering Education & Innovation and also the chair of the Engineering Education Faculty. She is also Director of the Sketch Recognition Lab and Professor in the Department of Computer Science & Engineering. She is a member of the Center for Population and Aging, the Center for Remote Health Technologies & Systems as well as the Institute for Data Science. Hammond is a PI for over 13 million in funded research, from NSF, DARPA, Google, Microsoft, and others. Hammond holds a Ph.D. in Computer Science and FTO (Finance Technology Option) from the Massachusetts Institute of Technology, and four degrees from Columbia University: an M.S in Anthropology, an M.S. in Computer Science, a B.A. in Mathematics, and a B.S. in Applied Mathematics and Physics. Hammond advised 17 UG theses, 29 MS theses, and 10 Ph.D. dissertations. Hammond is the 2020 recipient of the TEES Faculty Fellows Award and the 2011 recipient of the Charles H. Barclay, Jr. '45 Faculty Fellow Award. Hammond has been featured on the Discovery Channel and other news sources. Hammond is dedicated to diversity and equity, which is reflected in her publications, research, teaching, service, and mentoring. More at http://srl.tamu.edu and http://ieei.tamu.edu.

Tactile Learning: Making a Computer Vision Course Accessible through Touch-Based Interfaces

The term "visual learner" is a ubiquitous concept in education. It is often associated with experiential or example-based teaching that helps the students understand a concept through its application. However, for those with visual impairments, visual learning may not be an option. The Royal National Institute for the Blind reports that there are significant barriers to learning technical content for blind individuals, including access to visual resources and, correspondingly, difficulty interpreting visual concepts.

In this work, we discuss experiences and key takeaways from adapting aspects of a computer science graduate course to be more accessible for the blind. The course teaches advanced machine learning methods rooted in computer vision and gesture-based methods in order to build sketch recognition systems. Additionally, it emphasizes elements of human-computer interaction and interface design, many of which are visual concepts. In order to adapt the curriculum, we used a high-resolution tactile display capable of mirroring imagery from a video display into a depth map that could be felt. This enabled the dual presentation of visual content as tactile surface maps. Through this process, we learned several best practices in terms of how to create content that transfers well from one modality to another, and we also developed a number of guidelines for creation of teaching materials like notes and assignments in a way that is more screen-reader friendly.

This paper shares key takeaways while also communicating student and teacher perspectives on developing, teaching, and using more accessible materials. Our goal is to encourage instructors of technical courses that are traditionally visually-based to consider possible ways to enhance the accessibility of their curriculum.

Introduction

In 2013, the Royal National Institute for the Blind published a study highlighting several key issues for blind individuals in regards to access to information, including inaccessible technical notation and visual resources, as well as teaching methods that can rely too heavily on visual concepts [1]. These challenges can be especially difficult to overcome in highly technical fields like Science, Technology, Engineering, and Mathematics (STEM). While research has long indicated that blind students are interested in participating in STEM [2], recent analyses have shown there are still notable gaps in employment rates for blind or visually impaired individuals [3]. Involvement in academic research through the National Institute of Health (NIH) by individuals with a form of disability actually decreased from the years between 2008 and 2018 [4].

There are a myriad of factors making it difficult to address these challenges, including the ongoing braille literacy crisis that is key to developing certain technical capabilities [5]. Recent research has explored ways to use technology to help support braille literacy for technical fields [6,7], but that is only one aspect of the underlying issue. It is evident that the teaching methods and materials used in STEM courses must also be considered in terms of making the field more accessible, yet instructor awareness and education on this aspect of course development is generally low [8]. One potential positive out of the move to online classes brought on by the Covid-19 pandemic was the capability to explore new ways for shared and equal access [9]. However, even with the promise of technology-assisted access through online learning, teacher awareness and training remains a significant barrier [10].

In this work, we review the experiences of teaching a graduate-level computer science course with hybrid modalities. The class included extensive technical content in statistics, data science, and machine learning. Several units focused on the fields of computer vision and gesture recognition, so they featured many visual elements when demonstrating concepts. During lectures, we presented materials in-person with visual slides displayed on multiple large screens around the classroom. Concurrently, the lecture audio and slides were streamed through a virtual meeting room to remote participants. The slides followed careful design guidelines in order to be compatible with a tactile display for greater access. After a brief explanation of the design of the course, we discuss the approaches we explored for combining our traditional materials with hybrid learning tools and tactile displays.

Course Overview

Our class is a graduate computer science course that teaches advanced machine learning methods rooted in computer vision and gesture-based methods in order to build sketch recognition systems. Additionally, it emphasizes elements of human-computer interaction and interface design, many of which are visual concepts. In preparation for the course, we spent some time examining tools and potential techniques to improve accessibility.

First, we decided on details related to course materials and resources. The course materials needed to provide an equivalent experience for both students attending in-person and remotely. As the next section will explore, we found Google Slides to be the best option for presentation software with sufficient accessibility support. Slides and lecture notes were made available on Canvas, the default class management software for our university. Gradescope was used for quizzes and assignments as it had more accessible navigation than Canvas. Zoom was used as the virtual meeting platform as it was the standard for our university.

Next, we focused on finding a method to make visual examples accessible through tactile interactions. As we will discuss later in this work, creating understandable tactile graphics is non-intuitive for sighted instructors and is more involved than simply converting images into raised-line representations. We used the Graphiti Orbit tactile display [11], which is capable of converting graphic elements into a raised, tactile surface. The Graphiti has 5 levels of height, ranging from 0 (off) to 4 (max), and it can be set to different modes indicating color conversion settings. In the simplest mode, a grayscale image can be converted into 5 levels from white (off) to black (max), with intermediary levels corresponding to grays. It can also handle colors by converting to grayscale intensities with corresponding height levels. The resolution is comparatively low in terms of visual graphics, with a total of 2,400 pins in the 60x40 array, but this can work well with large images.

Finally, we experimented with different delivery methods for course content communication. Due to the hybrid nature of the course, our options included Canvas, Gradescope, Zoom, and email. We ultimately found that using a combination of different technologies worked best to provide easy access with high fidelity for technical content.

Course Design

Slides

With a preponderance of presentation software available today, it is common to see highly decorated slides. While this style of slides generally look very nice to sighted students, they ultimately contain a lot of visual noise that prevents them from being converted to a tactile representation using a tool like the Graphiti. We found that simple presentations consisting of a white background with stark black text worked best for everyone in class. Students benefited from the higher contrast and improved readability, especially towards the back of the room and over Zoom. There were few visual artifacts, and the minimal visual content did not pose as significant an issue for the tactile display.

In addition to making the slides clear and clean, we also wanted to ensure they were as accessible as possible to screen readers. Most presentation softwares, e.g., Microsoft PowerPoint and Apple Keynote, contain accessibility features such as screen reader mode and alternative text for images. Unfortunately, each of these tools requires slightly specialized commands to navigate the slides, creating potential issues. By contrast, Google Slides generates an HTML document when placed in accessibility mode. The benefit to using an HTML document is that it is highly structured and can be navigated like a standard website, making it familiar to blind or visually impaired students who use screen readers. Instead of using special commands to switch between slides, all of the content is presented in a single HTML document with headings for navigation.

Moreover, to facilitate easy listening of the entire set of slides, we added audio cues throughout the slide deck. We found that some screen reading software did not pause between bullet points, so we placed periods at the end of slide titles and bullet points to ensure that screen readers would read with a natural pause between listed items. To create an audio equivalent of slide numbers and gave a clear marker for the end of the slide, the final text on each slide read "end of slide \#." This final text was in a white font so that this informational detail was only accessible with a screen reader and did not add visual clutter, adhering to our aforementioned slide design philosophy.

Tactile Display

One of our biggest challenges was converting predominantly visual information into something that could be easily understood in another medium. The course includes many graphics, showing sketches and images in order to explain key machine learning features and algorithms related to core class concepts. Using the Graphiti display's mirror mode, we were able to take content presented over a full screen, share the screen on the Zoom call, and present it on the display. This meant that slides could be turned from visual to tactile using the raise dot system of the Graphiti. However, given the resolution limitation of the tactile display that is 60 x 40 pins, we had to construct images that could scale well to that smaller resolution. We found that by creating images with large thick lines we could retain key elements like edges from the original image onto the Graphiti. Through trial and error, we determined how various slide designs would convert to various pin heights on the Graphiti. We configured the Graphiti to put white as the 'off' height, creating a blank canvas. Full black raised the pins all the way, creating the primary tactile content. Through testing we found that specific shades of red and blue could be used as intermediary heights, allowing for additional details to be displayed tactically.

Another consideration for the Graphiti display is the refresh rate. In contrast to visual displays, the Graphiti cannot instantaneously change its display if a slide changes or updates. This quirk made standard animations indecipherable as the pins would constantly move and only output partial tactile graphics. To address this, we designed slides using a "storybook" or "flipbook" style, putting a single image on each slide when communicating examples that had multiple steps, e.g., a corner finding algorithm. If quickly navigated through, these images would create an animation of sorts. This approach gave students a high degree of control when reviewing the example, having the power to navigate the "animation" frame-by-frame by simply going through the slides.

Technical Materials

A third key component to our investigation regarded the delivery of technical materials, like mathematical notations and feedback on math-based quizzes and worksheets. While Canvas has some accessibility features, it did not prove to be easily navigable for highly technical material like mathematical equations in our case. We found Gradescope to be a delivery tool that enabled all forms of digital assessments, which had native support for LaTeX. That said, we determined the best notation practices for concepts such as matrices through some trial and error to maximize clarity.

While Gradescope has good accessibility support for taking assessments, it is still developing optimal ways to provide feedback. Feedback was presented separately from the assignment content, making it difficult to follow when using a screen reader. Instead, we use scheduled emails to send text versions of feedback.

Implementation

We had the good fortune of tackling the problem of tactile tutoring in the post-pandemic era using a Graphiti display and Zoom. With the help of an over-the-shoulder camera, the instructor could see where the learner's hands were and give verbal commands. As the lesson progressed, using an over-sized mouse cursor allowed the instructor to point to a particular part of the image. Because the Graphiti pins make some noise as they rise up and down, the result was a haptic and audible beacon that could be used instead of the clunky direction-giving that plagues most instruction of non-visual learners. Even something as simple as saying, "yes, that's it." is disruptive and takes the learner away from the context they are focusing upon.

Zoom has a screen share function which is much more useful than something like remote desktop control. With remote desktop control, there is the problem that you are actually controlling the mouse on the remote user's computer. While the image is the same, any attempt to use the mouse on the remote user's computer creates a cacophony of spoken screen reader descriptions as the mouse touches the different parts of the screen. Non-visual users generally navigate using the keyboard instead of the mouse; the mouse cursor is an object that the screen reader needs to announce. When someone else starts moving the focus around with a mouse, this makes it hard for the non-visual user to figure out where they are on their own computer.

We utilized multiple extended monitors to make the screen sharing fully effective. The non-sighted learner has full control of their usual desktop, and the extended monitor is used to transfer the shared screen to the Graphiti. This prevented any occlusion of the things that the learner needed to use to take notes or interact with their own computer.

An over-the-shoulder webcam on a tripod offered the instructor a zoomed-in view of the Graphiti surface at a significant elevation (7 feet up). This allowed the instructor to see where the learner's hands were situated while giving instruction. This almost overhead view of the display allowed the instructor to see everything in an upright orientation, with no personal space or contagion issues. This helped keep the lesson focused and free of any conflicts about personal space. One other useful part of a camera is that it could be positioned in front of the learner and the observer could rotate the view of the camera 180 degrees in order to see it from the learner's perspective. This would help avoid the problems of 'my left' versus 'your left' that plague direction-giving when people are face to face in an in-person setting.

Discussion and Insights

In this section, we discuss key experiences—and insights that may be learned from them—through both student and teacher perspectives. In general, while many of the practices we explored in the course design proved to be operational, there were still many takeaways in terms of techniques that could be further researched and improved.

Describing what part of the image the non-visual learner must touch to be in-sync with the instructor's actions

Verbal directions like, "find the upper-left side of the diagram, feel for the horizontal line and follow it until it intersects the diagonal line," take time to process, so are disruptive to understanding the content of the image. The non-visual learner has to interpret the verbal commands about where their hand should go, while also, and sometimes simultaneously, listen to the instruction regarding the content of the diagram. Because everything is flooding in via the sense of hearing, and the language is intermingled between directions toward and information about the particular subsection of the diagram, the learner has to parse these two streams of information separately, yet simultaneously. The result is a 'traffic jam' in the auditory cortex that visual learners do not have to disentangle. When a sighted learner sees the professor point to a particular part of the slide, they do not have to parse that separately from the words that are being spoken. The sense of sight smoothly integrates with the auditory input and linguistic understanding. This allows instruction to move at a rapid pace, but this pace often leaves the non-visual learner behind, or with missing pieces as the words often overlap simultaneously, causing information from both streams to be distorted or lost.

Making dynamic changes (AKA 'the delta')

Sighted instructors often use animations, multiple frames on a single page, or multiple slides to illustrate how one state is different from another. When an instructor uses two slides on a slide show with a subtle difference, the sighted learners will immediately see the difference between the slides, as their brain is cued to look for movement. When a line shifts on a background that remains the same, attention is automatically cued to the change. This barely requires any mention by the instructor in order to let the students know what has shifted, and that this must be the 'important part'. When a non-visual learner is given two slides with a single difference, they lack this convenient ability to instantly sense the thing that changed. It may require significant time to play the game of 'Where's Waldo' on two subtly different tactile graphics to identify the small difference between them. The differences do not 'jump out' from a tactile graphic the way they do for vision. An illustration would be to show two adjacent slides with a single change, but with several distractor slides intercalated between these. Most sighted people will have a much easier time with the two related slides played back-to-back because of the visual system's rapid attunement to the single change in a background. If the game of "Where's Waldo?" were played using sequential images rather than side-by-side images, it would be so easy that it would hardly merit playing. These subtle, yet powerful cues are used in teaching, but they are not nearly as useful in non-visual media.

Labeling graphics

When labeling images, many words can be jammed into the crevices of a visual diagram, but braille lettering is much less space-efficient. Most tactile maps resort to single letter labels and a legend to decode because of the space constraints that differ between braille and print lettering. This means that there must be a separate space or page for a legend, and that the learner has to shift from the label to the legend, often losing their thread of thought. In the world of the visual learner, this is akin to having to shift away from your immediate interest to a completely different area to 'look up' the thing you need. The journey is fraught, and the mental resources consumed by the seeking interfere with retaining the information. We've all had the experience of walking into a room and asking, "Why did I come here?" For a non-visual learner using a diagram with terse labels and a legend, much the same sort of thing happens, increasing frustration and cognitive load in a way that is not experienced by a sighted learner who may just look at a label and immediately tell what it is referring to in the diagram.

Multiple subject "blobs"

In the world of the sighted, multiple colors and shades are used to disambiguate the multiple subjects in an image. Subtle fill textures and color variations allow the visual learner to pull apart an image into the different parts. Imagine instead that you had just black and white at a fairly low resolution. Tactile graphics have just one ‘bit’ per taxcel (the tactile equivalent of a pixel). Most tactile displays will only have an ‘up’ or ‘down’ state, so the richness of a visual image is not available to support understanding in an image with multiple subjects. Only one refreshable display, the Graphiti, has multiple heights for its pins, providing 5 bits of information per taxcel with 5 discrete heights (0, 1, 2, 3, 4 and 5). This is not problematic for flat things like a map of the United States, where the lines that divide one state from the next are overlapping. It is much more difficult to keep overlapping subjects in an image from blending together, especially when they have not been pre-processed with tactile graphic ‘rules’ in mind. The cues used by the sighted to intuit that an object is in front of another object are not readily apparent to a non-sighted user. Visual artifacts like perspective and the ‘vanishing point’ have no tactile real-world equivalent. These concepts are solely artifacts of how vision works. When these subtle visual cues are translated into tactile graphics, the result is confusing, especially to those without visual experience.

“I can’t tell what you’re attending to”

The problems of perceiving another person’s position on a tactile graphic while sitting or standing in the same room are greater than you might expect. Humans are very adept at determining where another human’s gaze is oriented. The ‘whites’ of the eyes are a feature unique to humans among primates. Following someone’s gaze to see what they are looking at is part of the innate non-verbal communication that we employ. A non-visual learner may be disconcerting to the sighted tutor because they give no cues based upon eye position. It is very difficult to break the habit of looking at someone’s eyes to gain information, so in the process of a tutor learning how to teach a non-visual learner, there will be many missteps where the ‘eyes tell lies’. A tactile graphic reader may be exploring the graphic with both hands at the same time, so figuring out when to say, “You’re touching the right part,” may not be so simple. Indeed, saying such a phrase may itself be ill-timed (hands have already moved on) or disruptive. It is certainly difficult for a tutor to give verbal feedback in the midst of a lecture in an in-person class.

“I can’t read what you can read”

Printed braille is essentially a foreign language to the sighted. In spite of the simple encoding system and only 26 letters to master, reading braille by eye is not a skill that most people can or will pick up in the time that they work with a braille reader, unless they work at an institution that specifically serves the blind and low-vision. One of the most effective in-person techniques is to sit face-to-face across a desk, so that the learner’s body does not obscure the tactile graphic. This reduces the personal space invasion issue. However, adds another level of challenge for the tutor who has to read the image as well as the braille upside-down.

Non-visual learners are a diffuse, low-frequency group

There is little opportunity for any one person in any one physical location to become skilled at working with a blind student. At the university setting, a TA will have one experience teaching someone who is blind, and then will likely never get to use those skills again. With the rise of virtual tutoring and learning resources, we have an opportunity to hone and optimize the learning experiences of blind students by letting a virtual tutor become skilled in tactile graphic making and virtual tutoring. When geography is no longer a barrier, skilled ‘helpers’ can build a body of experiences with what works and what does not, in a way that was impossible in an in-person learning environment. Instead of clustering non-visual learners where the instructors know them as is common with schools for the blind and low vision, the instructors who want to do this work can ‘commute’ virtually, so that instructors at universities don’t all have to learn this skill set and then let it atrophy when that one student leaves their classroom.

Future Work

Author Amanda Lacy is currently working with undergraduate researchers to build a tool that would allow a virtual helper to send a portion of their screen’s pixels to the Graphiti in the exact proportions that it needs to utilize all of its display area [11]. This would be usable by the remote viewer to pan and zoom across a slide to highlight the important aspects of the image. In addition, they are building a ‘blinking dot beacon’ that can be placed by the remote tutor on the appropriate spot. Once the non-visual learner touches the dot, it is either dismissed automatically or by the non-visual learner who presses a hotkey. This will allow the tutor to silently place ‘look here’ beacons and concentrate on instructing rather than talking about where the information is.

Conclusion

In summary, our experience teaching this course has resulted in several recommendations we can provide to other instructors to improve accessibility for blind and visually impaired students. Naturally, course materials need to work well with screen readers. We found that Google Slides had the best accessibility mode by turning the slidedeck into an HTML document. The audio experience needs to be considered: adding periods to the end of bullet points and including non-visible end-of-slide markers improves the quality of the slide reading without adding visual clutter.

In our course, we presented lectures in a hybrid format with the slides screen mirrored onto a tactile display. This approach allowed the slides to be seen and touched simultaneously by the students. When going over examples, it is important to give adequate time to explore the tactile graphics, trying to minimize the amount of multitasking the learner needs to do to absorb the example and listen to a lecture. To communicate mathematical content, we utilized LaTeX notation to give the best screen reader experience. Our quiz and assignment platform, Gradescope, supported this option natively, facilitating the hybrid distribution of course materials.

When possible, materials should be communicated in plain text. We found that while there is some effort in creating materials that work well in different media and hybrid formats, many materials could be simplified to a text representation.

References

[1] Heather Cryer, H Cryer, Sarah Home, Sarah Morley Wilkins, Emma Cliffe, and Stacy Rowe. Teaching stem subjects to blind and partially sighted students: Literature review and resources. RNIM centre for accessible information, 2013.

[2] Bernhard Beck-Winchatz and Mark A Riccobono. Advancing participation of blind students in science, technology, engineering, and math. Advances in Space Research, 42(11):1855–1858, 2008.

[3] Michele C McDonnall and Zhen Sui. Employment and unemployment rates of people who are blind or visually impaired: Estimates from multiple sources. Journal of Visual Impairment & Blindness, 113(6):481–492, 2019

[4] Bonnielin K Swenor, Beatriz Munoz, and Lisa M Meeks. A decade of decline: Grant funding for researchers with disabilities 2008 to 2018. PloS one, 15(3):e0228686, 2020.

[5] Elizabeth R Hoskin, Morag K Coyne, Michael J White, Stephan CD Dobri, T Claire Davies, and Shane D Pinder. Effectiveness of technology for braille literacy education for children: a systematic review. Disability and Rehabilitation: Assistive Technology, 19(1):120–130, 2024.

[6] Laura Sauer. Mathematics for visually impaired students: increasing accessibility of mathematics resources with latex and nemeth mathspeak. 2020.

[7] Patric Galera Forcelini, Laura S´anchez Garc´ıa, and Ermelindo Paulo Breviglieri Schultz. Braille technology beyond the financial barriers: A braille literacy platform to effectively combat braille literacy crisis. In Proceedings of the 8th International Conference on Software Development and Technologies for Enhancing Accessibility and Fighting Info-exclusion, pages 41–46, 2018.

[8] Idalis Villanueva and Marialuisa Di Stefano. Narrative inquiry on the teaching of stem to blind high school students. Education sciences, 7(4):89, 2017.

[9] Amanda Lacy, Seth Polsley, Samantha Ray, and Tracy Hammond. A seat at the virtual table: Emergent inclusion in remote meetings. Proceedings of the ACM on Human-Computer Interaction, 6(CSCW2):1–20, 2022.

[10] Sathiyaprakash Ramdoss, Dongmei Liu, Shivaji Kumar, and Kaelin Lee. A qualitative study: Perceptions of students with blindness in post-graduate distance learning in stem fields. 2021.

[11] Orbit Research. Graffiti - a Breakthrough in Non-Visual Access to All Forms of Graphical Information. Accessed: 02/06/2024. https://www.orbitresearch.com/product/graphiti/

# # #