

Hello, everyone. I'm Shiri AzenKot, and today I'm going to speak to you about "Augmenting Reality for People with Low Vision." Let me get started by introducing the concept of low vision. This is a group of people with disabilities that I focus on in my research, and I also have low vision myself. So low vision is a visual impairment that can't be corrected with glasses or contact lenses. So it does affect someone's ability to live life, to perform daily activities and tasks. But, once again, it can't be fully corrected, so if you're nearsighted and you wear glasses, then you don't have low vision. And, actually, most people with visual impairments do have some vision-- they're not fully blind-- and yet this category, low vision, is largely ignored in the research literature and also in life in general. So it's a fairly invisible disability. Let me show you some rough visualizations of what low vision could look like. It can come in many forms. So let's say this photo right here of two boys holding two balls, let's say this is something a person with typical vision sees. So a person with low vision might see the image like this. So this is--again, it's a rough simulation of what someone with a cataract might see. So the colors are a bit dull, and it's a little blurry. And for a person with glaucoma, their vision is impacted differently, so they might have low vision. They might see the image like this, so their peripheral vision is affected. They would only see the center of the image. And for someone who has macular degeneration, which is yet another condition that can cause low vision, a very common condition especially for older adults, they might see the image like this. So macular degeneration affects your ability to see in the center of your visual field. So it's more difficult to see faces. So you can see that there are many different ways in which people's vision can be affected. And there are also different degrees. And it's actually very difficult to understand exactly what people can see. Perception is a very complicated thing. There are many tools that can help people with low vision. So, for example, this is just a handheld magnifier, and you can buy magnifiers at the local drugstore. You can also buy more powerful ones from companies that specifically cater to people with low vision. So, of course, this magnifies content that you hold up close. And people with low vision for decades have also been using various forms of electronic vision enhancement. So in this slide, you can see some examples. People use vision-enhancement systems that are handheld, and even head-mounted vision enhancement systems. So in my research, we are looking ahead at new forms of technologies and trying to understand how these new forms of technologies can be used to solve existing unsolved problems, specifically accessibility problems for people with low vision. So these days we see the increasing popularity of augmented-reality devices, and we are interested in using augmented reality as vision-enhancement systems. And specifically, we want to see how we can leverage augmented reality, how we can leverage this new technology to address visual challenges that can't be solved with these more traditional vision-enhancement tools that I just showed you in the previous slides. So let me distill this into the key research questions that I'm going to address here. So, specifically, a lot of my work focuses on this question: What augmented-reality visualizations will support people with low vision in daily tasks? And I'm going to talk about two projects that address this question. The two projects correspond to two daily tasks. Sorry. I was just adjusting all these annoying Zoom windows that pop up and cover my slides. OK, so the first task that I'm going to talk about is finding a product at a supermarket, a specific product, and the second task is navigating elevation changes, so, for example, stairs or curbs, any sort of change in the terrain that you are walking on. And before I go any further, I want to explicitly call out some wonderful collaborators that have contributed to this work. So Sarit Szpiro was a post-doc who worked with me. She's currently a professor at the University of Haifa in Israel. Yuhang Zhao was my Ph.D. student, and she is currently a professor at the University of Wisconsin, Madison. And Elizabeth Kupferstein is an optometrist, and she was consulting with our group and helping in all kinds of ways, especially in understanding visual perception. So let's

begin with understanding how we can address the first task that I mentioned, finding a product at the supermarket. And the first thing that we do in my research is understand the specific challenges that the user experiences in context. So we start out with a study. And for this study, we recruited 11 participants with low vision, our target population. And we had the participants arrive on our campus for a single session study. This study included an interview and an observation. We used a method called contextual inquiry, which is very core to my field, human-computer interaction. So in this method, we watch as participants complete a task, and we ask them questions throughout their performance of this task to understand exactly why they do what they do. So we gave them a task, and the task was this: Find a nearby pharmacy and purchase a specific Tylenol product. So the Tylenol--the product was Tylenol Extra Strength. Very specific. We told them to get, you know, the 500-milligram, 100-count box. And the idea here was that we were emulating a real-life scenario, where participants had to find a nearby business, especially a supermarket or pharmacy, and buy a specific product. In the case of a medication, the stakes are very high, so they did need to find-- they would need to find a specific product and be able to make sure that the exact details were correct. So the participants completed this task, and we observed them. They all found a nearby pharmacy, and, you know, this was done in a very dense area in Manhattan, so there were actually several pharmacies to choose from. They walked to the pharmacy, then they had to navigate within the pharmacy and find the specific Tylenol product. So they found the correct aisle, and then they were confronted with this scene right here, a bunch of different products that look very similar. And they had to find the specific one that we asked for. And this task right here, this particular subtask of finding the exact product among a set of similar products on the shelf, this subtask was the one that proved to be the most difficult out of all of the other subtasks within this larger task. So it was more difficult than navigating to the store and within the store, and purchasing the product and so on and so forth. And we found that there were really no tools that helped participants in this task because their job was not just to see the details, but they had to do a visual search task. So they had to scan the aisle and find the Tylenol product and then see the details to verify that it's the correct one. So this is very distinct than just reading fine print. So we decided to address this particular visual search task with a system that we call CueSee. So CueSee is an augmented-reality system, or a set of visualizations, and CueSee uses computer vision to recognize a target product on a shelf, so you can tell it which product you are looking for, and then it presents visual cues to direct the user's attention to the target. So in other words, it makes the target more visible to the person with low vision. And it augments the user's vision directly with augmented-reality glasses, so it superimposes the cues over the target product with augmented reality. So let me show you what this looks like. First, let's take a look at a mock grocery store aisle. So we set this up when we were going through the design and evaluation process of CueSee. So here are a bunch of products on a set of shelves, and it's actually really difficult to design visual cues for people with low vision because we don't know exactly what they can see. So the shelves might look a certain way to sighted people, but to people with low vision, as with the rough simulations that I showed earlier, they can be very distorted, and it's not entirely clear what visual cues will be most helpful at attracting people's attention and making a certain product more visible. So that was the design challenge that we encountered, and the way we addressed it was by turning to our knowledge of different types of visual conditions and also some of the literature from psychology about attracting people's attention. And we designed a set of five visual cues. So let me show you what they are. This cue is called guideline. So there's a red line from the center of the display to the product. This is another cue. It's called spotlight. So here we turn--everything but the product itself into grayscale. This cue is called flash. So we superimpose a border around the target product, and we change the color

from black to light gray. So we have this repeated onset effect that attracts attention. Then we have movement. Again, movement to something that attracts attention, especially with your peripheral vision. So the product here is rotating from side back to center to the other side. And then the final cue here is called sunrays. And here we have a set of lines from around the periphery, converging at the product. So by design, these cues can be combined. So here's an example of how two of the cues are combined, sunrays and spotlight. You may have also noticed that there is an additional enhancement here. The target product is magnified, and the contrast of the target product is also enhanced. So this is what we call the base enhancement. All the cues were superimposed on top of this enhancement. So we prototyped this system. This was a few years ago already. We used an Oculus. So we did this with what's called video see-through augmented reality, meaning people were seeing a video of their surroundings. And we were able to modify that video and add the enhancement and the cues. So we did this using the Oculus Dk2. And here's Yuhang Zhao, who prototyped this system and led this project. She's demonstrating here, using the prototype. So she's finding a product. You can see on the bottom right what she's seeing through the device. So the question was How effective were these cues, and, How well did they help people with low vision actually perform this visual search task for the target product? So we conducted a study to evaluate CueSee. And once again, we recruited participants with low vision. We recruited 12 to be exact, and there were two parts to the study. So the first thing that we did was assess the different cues. So we asked participants to try each of the cues and also the basic enhancements or the basic magnification. And that was our baseline. And we asked them to select their preferred combination of cues. Then using their preferred combination, we asked them to perform a product search task, so it was important to allow them to select their preferred combination because of the variety of different low-vision conditions. So we really wanted to make the system customizable and adjustable to the different degrees and types of low vision. So when performing the product search task, we had participants do this in two conditions, using their preferred cues and their best correction. So that means using glasses or whatever basic tools they used to correct their vision under normal circumstances. That was the baseline. Here are our key results. What we found is that all participants preferred anywhere between 1 to 3 of the cues, meaning that the basic magnification enhancement was not enough to help them actually find the product. They needed a cue to help draw their attention to the product and make the product more visible in the search task. And the different types of cues varied, so the preferences varied, but there were some common trends. For example, none of the participants liked the movement cue, and a lot of them liked the guideline and the flash cues. When looking at the comparison in the product search task, we found that CueSee was faster than the baseline, so that was a very good finding, and it was also more accurate. So there were no errors with CueSee, and there were a few errors with the baseline with best correction. So in other words, we found that CueSee was very effective for helping participants with a product search task. Let me move on now to the second project. So the second project looked at the task of navigating elevation changes, such as stairs and curves. So once again, the first thing we did was try to understand the particular challenges that the users experienced in context. So we conducted another study here, and we recruited 14 participants, again, with low vision. We always recruit target participants. We asked them to come to our campus, and we conducted a study where we--again, we interviewed them and then observed them complete the tasks. And in this case, the tasks involved navigating different areas around our campus that included elevation changes, like stairs and curbs. As they did this, we asked them questions to understand exactly why they did what they did, and we also observed them. Here are some examples of the places that they... navigated. So we asked them to walk up and down different sets of stairs. So the

stairs were in different rooms with different lighting conditions. Some had direct sunlight, some were in darker areas. The floors had different--were made of different materials, hardwood or concrete. So there was quite a variety here. And we also asked them to walk along areas that had other types of elevation changes, such as curbs or curb cuts, for example, when crossing a street, areas that had an uneven texture on the ground, and some areas that looked like there might be an elevation change, but really there was none. It was flat. What we found was that three of our participants used a white cane when navigating throughout the study. And this was interesting because, again, you know, typically, we imagine people with visual impairments to navigate with a mobility aid, but it turns out that many do not. 11 did not use any tools, although, interestingly, six had a cane with them. They just chose not to use it. We found that walking up and down stairs was the most challenging of the subtasks that participants encountered. So in other words, more challenging than walking up-- than walking along curbs or curb cuts or any of those other examples that I showed. And participants walked very slowly. They paused a lot. They looked down, they used other senses. They shuffled their feet a lot to feel the ground. And they also touched wherever they could. So in the case of the stairs, they touched the railing for extra support and indication for what was going on. So to address the challenge of navigating stairs, we designed a system that we call StairLight. And StairLight is, once again, an augmented-reality system, and it recognizes stairs with computer vision. So same idea, so we're taking advantage of computer vision here to compensate for some of the perception difficulties that the user encounters. And then we present visual, and, in this case, also some audio cues to enhance the salient information for the user. So in this case, that information is the edges of the stairs. So we enhanced the visibility of the stairs with augmented reality. So with CueSee, we used a head-mounted display, but with StairLight, we decided to use projection-based augmented reality. So this is something that's looking a little farther into the future, and we see trends that promise to have projectors, augmented-reality projectors, more commonplace in the future. For example, present on cell phones and mobile devices. But for this particular prototype and evaluation, we used stationary projectors just to evaluate our design and to get an idea of what this could be like. So here's an example of the projector that we set up in the environment and an image of what it looked like in our prototype. So here's someone walking up the stairs, and she is using the augmented-reality projections... to help her see the edges of the stairs. And just like with CueSee, we experimented with different cues, cue designs, to accommodate for different low-vision conditions. We did the same here. We experimented with different highlights for the stair edges. So the most important edge to highlight was the top stair. This marked the beginning of the staircase. So we looked at different ways of highlighting it. For example, here's a flash. We also used stripes and motion, so horizontal motion and also vertical motion. And for the rest of the stairs, we looked at two different visualizations, blue and yellow. And these were designed to be less visible than the top stair because we wanted the top-stair visualization to really pop out and attract people's attention. So as with CueSee, once we designed StairLight, we wanted to evaluate it with participants with low vision. So we recruited 12 participants with low vision. And we did two things once again. So the first was we assessed the different visualization options. So we had participants try each of the visualizations that I showed and select their preferred combination for the top stair and the intermediate stairs. And then we had them do stair-walking tasks, so walk up and down stairs. And, once again, we had them use two conditions, so their preferred visualizations with StairLight and then the baseline, and the baseline was their typical walking method, so if they used a cane, we had them use a cane, or if they did not use a cane or any other mobility aid, we had them do that. And what we found was that most participants-- seven out of 12--preferred the basic visualizations, so they didn't like any of

the motion animations. They thought it was distracting. And then the rest of the participants-- so a little less than half-- had varied opinions. So there was some common trend there. The motion was a little distracting for most. We found that with StairLight, the walking time was reduced... in most cases, so that was a good find. There was a strong trend there. And the subjective experience that people reported after experiencing the system was positive. So overall, we had positive findings for this system, as well. Those were the two tasks that I wanted to focus on here, two ways in which we used augmented reality to enhance the experience of people with low vision in two specific daily tasks. And the idea more generally in my research is to look at how people perform, specifically people with disabilities, perform daily tasks, whether it's navigation, whether it's grocery shopping, whether it's working, reading, socializing, surfing the web, so any sort of daily task to understand the challenges and to try and leverage emerging technologies, in particular augmented and virtual reality, to help them be more effective and productive and enjoy what it is that they're doing. So moving forward, there are many interesting questions here to think about, some lower level, some higher level. I'll highlight a few. So customization is always an interesting question because there's such a range of visual conditions, and, in general, with people with disabilities, there is such a vast range of not just the degree of ability but also people's preferences, their prior experiences and their desires. So it's important to consider those things and to design systems that can be customized for the user without being overwhelming. Sensory output. I think it's important to try and understand when we should use what sensory-output visualizations, audio output, and so on and so forth. And perception and behavior. The next step is to understand how these technologies affect people's perception and their behavior in real life, so to close the loop. Ideally, when this technology becomes more mainstream, I'd like to work on incorporating our work into consumer devices to have a real-world impact. But right now, we are faced with the problem that all of the augmented- and virtual-reality technologies that are coming out on to the market are incredibly inaccessible to people with low vision and to people with other disabilities. So for this reason, I co-founded the XR Access Initiative, and this is something that I wanted to highlight here. So the idea behind the XR Access Initiative is to bring together academics and industry leaders and also advocates, people with disabilities, and anyone else who is interested and engaged in this space, and to make sure that we are from day one--day zero, even-- we are thinking and working on making these emerging platforms accessible, so instead of waiting until they are already popular and widely used, to think about accessibility, we need to do that now. Before they are commonplace in schools and in the workplace, we have to make sure that as they are being designed, they are accessible to as many people as possible. So XR Access has been in existence for several years now, and we have an annual symposium every year. This year, it's coming in June, and we are currently trying to figure out whether it will be hybrid or in-person or virtual. So check out the website for the latest: xraccess.org. We also have research talks. They are once a month. So if you join the mailing list, you'll get updates on those. We also have a research experience for undergraduates on the topic, so tell your undergraduate students to sign up. The deadline is March 15th. Again, go to the website, xraccess.org. So lots of exciting things ahead. And I will just finish up by thanking the National Science Foundation for funding a lot of this work and also Verizon Media, Facebook, and Google for partially funding the work as well. I'd also like to thank my collaborators and everyone here for listening to this presentation. Thank you very much.