My name is Nicholas Giudice. I'm a professor at the University of Maine, and I'm excited to be speaking with you in this Conversation Series. So, this talk is about the role of information access and scientific visualizations for STEM disciplines and, really, for any fields requiring scientific and computational thinking, understanding, and processing. And my goal today is to consider limitations of current visualization approaches and to discuss best practices for improving how we create and visualize data, really, of all kinds so it increases general learning and understanding, so it increases inclusion for those who consume it, and so it better mirrors how the brain actually works and processes information for visualization. And note that while this-- there's a lot of science underlying what I'm talking about today. Given the time constraints, I'm not focusing on data, but what I really am concentrating on are high-level concepts that I hope will generate some subsequent conversations and feedback. All right, a little background about me. So, as a blind guy myself who got my Ph.D. in a vision science program and who now works in a spatial computing department focusing on multisensory information access, you know, I bring a somewhat different perspective to the normal view on the role of visual information and visual perception and the technology that we use to visualize data. So, the traditional approach for conveying numeric data, whether--when we're communicating ideas, and whether they be abstract or concrete, involves using different visualization techniques. You know, we have diagrams, figures, animations, and what have you. A major theme of this talk is that most of these approaches are outdated in terms of ideology, in terms of an understanding of human information processing, and in terms of the technology used, and we'll discuss specifics of each, but the aggregate result is that much of what exists and is used in STEM fields today is exclusionary to many people and does not maximize the brain's native abilities, and that's something that we really want to do more. That's a big kind of take-home for me, so the alternative approach that I'm advancing here argues that dynamic-- any type of numeric and graphical representations and the technology that we use to produce them should be designed based on how the brain actually works and how it uses multisensory information and multisensory interactions, and the result of using this bio-inspired, multisensory approach is that our visualizations will be more intuitive, they'll be more realistic, and they'll be more inclusive. A second, quick note-- although a lot of my--the work that I'm talking about today is geared toward, you know, improving access, specifically for blind and visually impaired people-- this is a focus of a lot of my own research; it's also a personal interest-- inclusive design is, by definition, independent of any ability or disability, so as such, the ideas and best practices that I'm touching on here are relevant to improving access, inclusion, and equity for everybody, and that's really the goal of universal design, more broadly. OK, so I've structured this talk around a few problem statements that I argue underlie a lot of the issues that I'm discussing, and then I'm following these up with a few positions and hypotheses that not only make my point, but also serve as some guidance for moving us toward needed solutions. So, my first problem statement is that most of the tools and technologies and terminology that we use to perform or talk about visualizations are based on a misconception about the visualizing process. And so, you know, whether we're talking about data visualization or visualization of a process or design, conceptual visualization, visualizing something in our heads, what have you-- there's tons of types of visualizations-- they all make a mistake, and the mistake is in the belief that... what is happening is about vision and visual information and, you know, this is kind of an understandable confusion here. You can't write, you can't even say "visualization" without including "vision" in the process, but I'd argue this is a little bit too literal, so this leads us to my first two positions. All right, Position 1. This is characterized by the following hypothesis. The majority of what most people call visual information isn't really visual, but is actually spatial information, and I'm not talking about outer space, but I'm talking about the spatial relations of the world around you. And a lot of people push back on this immediately. It seems kind of not intuitive, but I invite you to look around the room that you're in right now, and my question is what--as you're looking around, what is intrinsically visual in what you are seeing, all right? So color, certainly, is fine. Texture is a good example, but think about the 3D structure of space: the relations between edges and surfaces, the direction and distance between objects and people. These are things that you're seeing in their spatial properties. They can be specified through vision, and that's a wonderful way to do it, but also through other spatial channels; touch and haptics, audition in hearing, and spatial language being the most notable. OK, so my second position kind of builds directly on the first, but it's tied specifically to visualization. So P2 argues that since the vast majority of visualizations rely on vision and visual information, and since most visual information is actually spatial information, we get the logical conclusion that most visualizations actually convey spatial information. So again, think about some of the types of representations and numeric representations that you use, whether they be data graphs or figures or diagrams, photos, maps, 3D renderings, what have you. As with the room example, these all mostly are about representation of spatial information, spatial relation, and spatial configuration and, again, are best described and are best depicted using spatial concepts, rather than visual concepts. So the previous two positions motivate a second core problem for information visualization. So problem two is that most visualizations rely only on the visual channel for conveying information. This is kind of expected from what I'm saying, but this leads us to a third position, and that is, essentially, since the brain processes spatial information from all of our senses, as we just described, effective visualizations should be multisensory, so, for example, they should use multiple channels of information for delivery, not just vision. I'll discuss more on specific benefits of multisensory... you know, the advantages of multisensory information, but the bottom line is that while vision is a great conduit of spatial information, it's really this, you know, really broad, huge pipe to the brain. It doesn't have a monopoly on space. All of our senses convey spatial information and so, therefore, I think of space as what you can think of as the common denominator of the senses, and I know I'm kind of a skipping record on this point, but I think it's hugely important for dealing with the issues that we're talking about today. And one of my key take-homes is that we-- and I'm talking here as the collective "we"-- we need to do a better job of appreciating this point and considering how to use multisensory information when we design technology and techniques for supporting visualizations of all kinds. I want to talk very briefly about some underlying theories. I think this is important if you were to accept my arguments here of why we need to use multisensory visualizations and why they work. So there's now results from a growing body of converging research-- behavioral studies and our imaging studies, computational studies-- and they're all demonstrating that there is far more similarity than difference in the processing computation and representation of multisensory spatial information in the brain. I've worked in this area for longer than I like to think, and I've done a lot of studies at this point, kind of looking at this idea with a number of different colleagues, and this is based on what's called the functional equivalence hypothesis of spatial information. Kind of a mouthful. There's a lot of underlying stuff here, but the essence of the argument is that spatial information learned from both... from different input modalities, whatever they are, whether they be touch, vision, audition, linguistic information, however you learn it. This leads to spatial understanding and spatial behaviors that are statistically identical, so functionally equivalent, independent of what sense that you use during the learning process. I think maybe this example might make this a little clearer. So imagine that you're learning this simple route map that we're showing here, a 3-legged route map, and you're gonna learn it from vision and you're gonna-- or learn it from touch, OK? If you develop functional equivalent representations, then you should be able to do things like pointing between remembered objects or locations, re-creating the map, or, you know, various other behaviors with similar accuracy independent of how you actually learn the map, and it turns out this is the case. We've done this with lots of studies. I'm not showing the data; if you're interested, these references show you a number of some of the basic work... but this brings us to our fourth position, and Position 4 argues that the robust evidence supporting functional equivalence that we just talked about is best explained by what we call the amodal hypothesis. And the bottom line of what this means is that spatial information from the different inputs that we talked about that you might use during learning, this all builds up into a single amodal or what you can think of as a sensory-independent spatial representation in the brain. And this... this representation is not tied to the input source or the input learning modality. So, therefore, functional equivalence occurs because our computational and our cognitive processes and the behaviors they support are tapping this unified representation when they're guiding our interpretations and our actions, and we call this amodal to 3D working memory representation the spatial image. And again, this is just a block diagram showing you kind of conceptual framework of how it would, you know, it would be implemented. OK, so, at this point, you may be saying, "All right, Nick, that sounds reasonable," or hopefully you are, at least, "but, you know, what's the connection here "between functional equivalence and amodal representations and visualization, which is how you started?" So the bottom line is that this basic research speaks to the ability of non-visual information to support the same level of performance as can be done using visual information, and this is important because it provides a theoretical framework for why communicating spatial information using multimodal approaches works for scientific visualizations. And there are many advantages of using these multimodal interfaces, but my-- I'd say my fundamental take-home is that our interface design, especially for visualization, should borrow more from what our brain does so well. Our brain is amazing at seamlessly integrating multiple channels of information into a unified experience. Here are some more practical kind of benefits of using these multisensory interfaces and types of visualization. So, first, doing this optimizes the brain's ability to interpret, process, and act on spatial information. So, in other words, you know, for a visualization to be maximally efficient in how we-- in how we use it and what it's communicating, we should complement the traditional visual content with multiple sources of non-visual information. Second guideline-- designing visualizations to be inherently multimodal and using redundant channels of spatial information, so, again, using different ways and different senses to provide redundancy represents inclusive design that's going to benefit the greatest number of people, so it's gonna be beneficial to people that have-- with different disabilities, people of different ages, people of different learning styles, and, I'd argue, really everybody. We are all multimodal learners, you know, otherwise we would have evolved with just vision, and that's-- that's not the case. Third, and perhaps the most obvious, is that using multimodal visualizations is most consistent with how we actually perceive, experience, and interact with the world around us, so we do this based on this rich mosaic of multisensory information, and why shouldn't our visualizations be doing this? So, why is that? Why are rendering techniques and user interfaces still based almost exclusively on vision and visual information and largely ignoring what is-- what we've learned from, you know, types of studies that we've talked about and, more importantly, from the most important information processor in the world-- our brain, right? The brain is better than any computer, even a tiny mouse brain, for simultaneously perceiving, processing, representing, and acting on information from multiple senses. So, the reality, however, is that the vast majority of technology that we use to visualize information is still limited by simple interactions, like we do typing and we have tapping and we have finger gestures for input and data manipulation. We have visual displays for output and information presentation, and this is silly. It's not how we actually experience the world through all of our senses. There's lots of reasons why, right? So, partly, it's because visual interfaces provide the most bang for the buck. This is something that's easy to do, people are used to vision, and it's what has been developed over the years. It's also partly due to inertia, so, you know, people-- it's what's worked good enough in the past, and there's kind of a basic human nature. People resist doing new things if what they're currently doing is working good enough. This brings me to my final position, Position 5, and it argues that good enough doesn't cut it anymore. Our visualizations need to adopt new bio-inspired approaches for conveying information based on how the brain uses multisensory input and output, the things that we've talked about, and this is frequently talked about, a lot of people will say this, and there are many known benefits. We've talked about some, there are many others, but the reality is that there is little actual progress in what's been done and what's--and, you know, how these visualization technologies have advanced. And this, as I was writing this, made me think of my grandmother. My grandmother used to tell me when I was equivocating or not doing things, and she'd say, "Nicky! Do your business or get off the pot!" I think she would not like my representation of her voice. Either way, it's a good point. I'm listening, Gram. That's why I'm trying to give this type of talk and get the word out. There are some really promising tools in the visualization domain that are doing the types of things that I'm talking about, and specifically augmented reality and virtual reality. There's lots of variants here. You can do this in lots of different ways. The technology can use monitors, it can use head-mounted displays, it can use caves, you could have AR glasses, but the advantages in general of this technology for visualizations are that they're based on 3D simulations, they're highly immersive, they allow for 3D-- 3-dimensional user movement and interactions, and it supports the ability to model and simulate any type of multi-dimensional data. That's a really big deal, and I'm particularly excited about this technology because it finally is moving from vision-only interfaces to using multimodal information, and this is, you know, this is important because historically, virtual reality was synonymous with visual reality and visual simulations. If you were in a VR world, that's what you got--visual stuff, but this is now changing and, you know, so--for instance, we're getting VR systems that are starting to use spatialized audio, so this is where you hear sound that's coming from around you in a 3D space, they're using touch and haptics, they're using temperature or virtual temperature changes. They're even using taste and smell in these simulations, and so this is--this is important, it's beneficial. It means that, you know, by using these cues, you're not only going to increase VR inclusion to people who can't see and otherwise couldn't use it, but you're also going to greatly improve the realism and the impact for everybody because we are now finally modeling how the brain actually receives and processes information in these multimodal interfaces. And the bottom line here is that both VR and AR have been used for some really amazing visualizations in lots of different domains, and the growing focus on moving beyond the kind of vision-only interfaces, I think, is hugely promising for visualizations of the future. This is something that I think is going to be particularly important. OK, I'm going to end by quickly discussing one line of research in my lab that's based on multimodal, bio-inspired visualizations, and there's a lot of projects that I'd like to talk about, but I have time for one, and this one I'm doing because I think it's particularly important. So, currently, there are over 12 million people in the United States alone that have some form of uncorrected vision loss, and this balloons to 280 million people worldwide, so we're not talking about a tiny demographic, but--and most of these people have significant trouble accessing visual graphics because there's currently no easy way to produce or convey graphical content non-visually. So our goal was to say, "OK, how can we develop new multimodal visualization" techniques based on "a lot of things that we're talking about that can be used in STEM fields of all types?" And so our solution uses the touch screen of smart devices, so phones and tablets to communicate graphical and numeric representations through combinations of touch, auditory information, high-contrast vision, and kinesthetic cues. So you may say, "OK, well, that sounds interesting, "but, you know, a touch screen is a flat glass surface," there's no tactile cues. How are you actually using touch on this?" What we realized is that the embedded vibration motors that are in phones, for instance, are actually amazingly good at generating tactile sensations. This is called vibro-tactile cue, so we're used to using this stuff as tertiary cues. You know, your phone rings and there's a vibration when there's an alarm, but you're not using it for anything that's really substantive, but we can use this as a primary cue as well. So, how it works is that whenever you have an onscreen visual element-- it could be a map, a figure, what have you-- whenever it is touched by the finger, as someone's moving their finger around the screen, we can trigger predefined vibration patterns synchronously at that screen location, and so the result is that, as people are moving their finger around and feeling the lines or the points, they're able to feel these vibrating surfaces and vibrating elements. And I said--you know, as I said, they can be points, they can be lines, they can be regions, and they can be traced. It's really quite compelling, so...you know, where the-- where the traditional paper-based technol-- tactile renderings--and this is kind of the current gold standard for haptic visualizations, but they're static, they're touch-only, and so they have a lot of limitations. Our system here is dynamic, it's intrinsically multimodal, as the haptic-based spatial information that I was talking about feeling can be combined with auditory cues to provide semantic information, they can be provided with--using enhanced visual information 'cause a lot of people with low vision still have some vision, and so we end up with a really robust, multisensory visual experience, and we call this the vibro-audio interface, or the VI. I've now done-- we've worked on this for a bunch of years, and I've done a lot of research in the lab using this visualization tool, and the results are hugely promising. So some of our studies are evaluating its efficacy for learning things like maps and figures and different scientific renderings. Others are looking at common data representations that you might use in school or stats or vocational settings like bar and line graphs, pie charts, Venn diagrams, and the like. And the reason I am so excited about this approach is that every study we've done, and we've now run hundreds of participants, has shown that the VI system that we've developed is equivalent to or better than current gold-standard approaches for non-visual rendering. But unlike these approaches, as I mentioned earlier, the VI has the benefit of being dynamic, so we can update and things change as people are feeling them. It's multisensory, it's portable, and it's implemented on smartphone, which is a computational platform that 80% to 90% of blind people already own, so there's a huge benefit for penetration and actually getting out there to be something that people can use. Really, of all the work that I've done in the last 20 years on accessibility, I'd say that I'm most excited about the promise and the prospect of this technology for improving information access and multisensory visualizations. Another thing I'm excited about is we've consolidated a lot of this work into several published guidelines papers, and these are giving design guidelines for how you'd use this interface. I think this will be hugely beneficial for people who are interested in using the VI in their visualizations and their curricula moving forward, and people are starting to do this around the world, which I think is really great. All right, I'm going to finish by summarizing a few take-home points that you may be able to predict from what I've said, so first, you know, we need more bio-inspired visualization approaches based on the brain's playbook for natural information-processing. Second, we need to move beyond the vision-only bottleneck of visualization to support use of multisensory information for both input and output channels and how we present and communicate information. Third, we need to design our visualizations using technologies and techniques that serve as many people as possible, so this inclusive design, I think, is critically important, but it also needs to be done from the onset, instead of a post hoc Band-Aid. A lot of times, people try to make things inclusive at the end, and therefore it really doesn't work. Fourth, and finally, we need to-- we need a lot more cross-talk between scientists and engineers and tech developers, the people that are making and using visualizations, and the end users that are, at the end of the day, trying to consume these to be useful. I think this will lead to far better products, far better techniques for supporting inclusive visualization of the future, and I think that this is absolutely possible. I want to thank you. I look forward to input and feedback that you have, and I will look forward to hearing from you.