Thanks for joining me today. My name is Adriana Arcia, and I am a nurse scientist. I've recently joined the University of San Diego Hahn School of Nursing and Health Science after 10 years at Columbia University School of Nursing. My faculty role means that I conduct research as well as teach, primarily research methods. Today I am going to talk about what visualization lessons we can learn from the needs of lay audiences. The work you're going to see today isn't solely mine. It represents the efforts of many people, including co-investigators, project managers, and professional designers. I could talk about this topic for days, but I have tried to keep this talk on the shorter side to leave time for questions and discussion.

Before I launch into my talk, though, I want to make a plug for a call for papers. This will be a special focus issue of JAMIA, the journal of the American Medical Informatics Association, about visualization of health data for lay audiences. The QR code will take you directly to the call. I am one of the four associate guest editors for this issue, so you're welcome to reach out to me if, after reading the call, you have questions about the scope. I will show this information again at the end of the talk, too.

For 10 years now, <click> the focus of my research has been on using information visualizations, like infographics, to support the health information needs of people with low health literacy and limited English proficiency. My team uses participatory design methods for our projects, so that means we're always working with members of the target audience to ensure that we are designing visualizations that are easily comprehensible, culturally congruent, actionable, and visually appealing. The visualizations are intended in some cases to return research data to participants, and in other cases to communicate clinical data to patients. I'm showing you English versions, but everything also has a Spanish version as well. I'd like to start by defining what I mean by the terms you see in the title to this slide, starting with information visualization …

…which I want to differentiate from data visualization. As members of DVS, you are probably focused on the "explore" side of the explore/explain continuum. Based on what I see in the Slack channel, a lot of you are involved in visual analytics or making dashboards. Information visualization is on the "explain" side of the spectrum. It's quite possible that your work has more aspects of infovis to it than you think. In infovis, the data have already been processed to some degree to move us up the DIKW pyramid. Information visualizations generally have a message and a point of view. Even if you would judge the information being presented to be very neutral, it still has an editorial slant because someone has made decisions about which information to present and what to say about it. Now that we're on the same page about the term information visualization, let's talk about tailoring.

Health care has a very long history of using visual aids for clinical care. We've all seen anatomical posters, food pyramid recommendations, instructions for how to do physical therapy exercises, etc. However, the vast majority of those visualizations I would classify as generic, which is to say that they are the same for every viewer; everybody sees the same poster. When we tailor a visualization, we'll actually modify the image in some way, usually based upon data we have on the viewer. In the case of the example at right, we have tailored this infographic with the prolonged stress score of the viewer and are comparing it to others in their cohort. Of course, we don't have artists stationed in every medical exam room, so tailoring is made possible at scale because of technological tools that can automate the tailoring for us. Some of what I do is work with programmers to create the bespoke software that tailors the visualizations that we have designed.

As an aside, I'd like to further differentiate tailoring from personalization, which is simply the practice of putting the viewer's or reader's name on something. A lot of your junk mail will be personalized, but not tailored. Lastly, I want to define what I mean by lay audiences.

For my purposes, I am defining lay audiences as people interacting with visualizations in a non-professional capacity. In the realm of health visualizations, that means patients, caregivers, community members, research participants, etc. The individual in question could potentially be a healthcare professional, say a surgeon, but if that surgeon is being shown a visualization by their own primary care provider, in that instance they would be a lay audience member. I suspect that most of you do work geared primarily toward professional audiences. That is, people who are using the visualizations you create as part of their job. So, you may be wondering about the relevance of this discussion of lay audiences and their needs to your work. I hope that by the end of my talk, you'll appreciate that the needs of lay audiences and those of professionals may not differ hugely, and therefore we can pull out some lessons that can make visualizations easier for everyone to interpret.

So, what are some of the key differences between lay and professional audiences? <click> Lay audiences are likely to have a broader range of educational attainment compared to professionals, and in many cases, the average level of education will be lower. <click> It's not a 1-to-1 correlation with education, but literacy, numeracy, and graph literacy are likely to be lower, too. <click> Regardless of educational attainment, professionals are – almost by definition – going to have a lot more domain knowledge about the subject matter of the visualization than lay people will. That means that you can expect professionals to have a certain amount of base knowledge upon which you can build and which they can use to interpret and contextualize whatever you're showing them. That *can* allow you to take some shortcuts and get right to the point or leave some things out because those bits are simply assumed. <click> One of the most important differences is that if someone is using a visualization or a dashboard or what have you as part of their job, it kinda doesn't matter if the usability of what they're looking at is poor. They have to

persist anyway because it's their job. They'll eventually figure it out, or learn the requisite workarounds simply because they have to. Lay audiences may or may not have the same level of motivation to persist but even if they do, their problem-solving skills might not be as robust or they may not have a co-worker or manager to turn to for help when something doesn't make sense. There may not be many human or technological supports in place. <click> As a result, for something to really work for lay audiences, the quality has to be very high. <click> If a professional sees a visualization and there's a typo or some information is left out, they can often compensate, "Oh, the topic is GDP so they probably mean billions of dollars, not thousands." Lay audiences may not be able to compensate and fill in the blanks in a similar manner. So, anything you make for lay people just has to be better. There's very little room for error.

What I'd like to do now is give you a bit of a recap of what I've learned over the last 10 years of working with lay audiences. The way this all got started was that my mentor, Dr. Sue Bakken, was the principal investigator of a research project at Columbia University that included a survey of nearly 6,000 residents of the local community. She felt that we have an ethical obligation, as researchers, to not just grab the data and run, but rather to return the data to the participants and the local community in a way that they would find useful, easily comprehensible, and actionable. Her idea was that we could use infographics to do this. We quickly discovered though, that there was very little in the scientific literature to guide us in creating these infographics. There was some research about how to convey risk, for example, what proportion of patients will experience a side effect with medication A vs medication B. But nothing about how to show someone their blood pressure, for example. So the focus of our earliest work was on finding the optimal graphical formats to display different types of data. At that time, I tried to use as little text as possible because I wanted to minimize the reading level but also– to the extent possible with qualitative methods -- to really isolate out the effect of graphical format.

Here are 4 of the prototype designs we tried out for showing blood pressure. All the designs show the same data, but in different ways. I'd like for you to take a moment to think about which design is easiest for you to understand. Now, I want you to think about Maria. Let's say she's a sixty-something year-old woman who was born in the Dominican Republic and left school after 8th grade. Maria came to the US as an adult and speaks minimal English. Further imagine that she's seeing the Spanish versions of these infographics. Think about which one Maria would prefer and write it in the chat: A, B, C, or D.
<click>
If you chose D, you're right! If we had just followed conventional wisdom that "less is more" and to "keep it simple" for lower literacy audiences, we would have gone with the stoplights. But as it turned out, D was the most popular because it was the most informative. Let's take a closer look at the final product after the graphic designer cleaned it up.

Here's what makes this image successful. 1) It is information-rich. Just because people might have lower literacy, that does not diminish their desire for information. It just has to be presented in an accessible way. 2) It uses a familiar color analogy. Stoplight colors are very powerful in their ability to communicate. We do need to be thoughtful about our labeling and the specific shades of red and green that we use so that we don't disadvantage viewers with color blindness. 3) The number lines support comparison. We can easily compare this patient's blood pressure numbers to the reference ranges to make sense of them. And 4) it provides context. The image on the upper left reminds the viewer how we arrived at these values and the figure on the right tells us why we should care.

What I learned about graphical formats is that for many of the study variables we wanted to communicate, reference range number lines were a big winner. Here's another example of a reference range number line supplemented with body silhouettes to support interpretation.

Another format that worked well was the visual analogy. Anyone who has ever owned a cell phone understands the consequences of getting into the red like the icon on the right. But it's almost more instructive to look at what didn't work.

One of the things we discovered is that people are very literal. Like super, super literal. If you don't read Spanish, the title says Fruit Servings Per Week. Here we're comparing how many servings Victor ate to what was reported by other men his age and to the dietary recommendation. When I showed this to my participants, they only talked about apples. "Wow! Victor really likes apples!" "So many apples!" "Victor eats apples every day!" They didn't generalize to other fruits. Occasionally, someone would do the mental math and point out that the recommendation only worked out to two servings per day, but most people didn't do that. <click> I had been avoiding showing servings per day because I wanted to use only whole numbers instead of decimals, but it seems avoiding mental math is more important. More importantly, I wanted people to understand that we meant fruit generally, not just apples. So, the participants told me that to accomplish that, I had to show a *variety* of fruits.

So I dutifully created these icons combining a variety of fruits. The people who saw the apples and then saw this said, "Ah! This is much better. You can see that there are a variety of fruits!" But I tested again and only showed this without showing the apples. You know what they said? "Victor eats the same fruits every day!" I had my sweet little old ladies carefully studying the icons, "banana… apple… grapes… a WHOLE PINEAPPLE!" Super literal.

Here are a couple more examples of literal interpretation. We had participants tell us that Victor has 4 pairs of shoes whereas the other men have 3. They told us that these are the kinds of shoes you should wear to exercise. They told us that Maria was active. That she was fast. They even told me that she was an educator and did good things for her community but many participants could not tell me how many days a week Maria exercised. Furthermore, the conversation was about walking, running, and jogging, but participants never talked about other forms of exercise like swimming or soccer. And they couldn't give me any suggestions for what image would convey exercise more generally better than the jogger. What we figured out was that in the presence of this propensity for literal interpretation, it is very difficult to represent concepts that are low-level abstractions like fruit or exercise. Furthermore, this graphical format was simply not familiar to our participants.

Which is really curious, right, because the stars were a big hit and worked really well. Why did the joggers fail? They are fundamentally the same format, no? Two reasons. 1) Our culture has trained us how to use the stars. Whether it's online shopping or movie ratings, the stars, specifically, are a familiar convention. I took Edward Tufte's one-day workshop back in 2013 or so and I asked for his advice about communicating effectively with my audience and he said that I should pay attention to what works in the wild. The effectiveness of the stars is the kind of thing he meant. He also implied that I was underestimating the abilities of my audience, and on that point he was very sorely wrong. I once worked with a participant who couldn't interpret the number 1.4. She didn't understand simple decimals. And although most of our participants could interpret simple bar charts, which is what we ended up using for fruit and vegetable servings, that was not the case across the board. Anyway, that's reason number one for the differences between these images: our culture has trained us on the star rating system. Reason 2 is that the image on the left is clearly metaphorical and represents degrees of something. People don't think we're telling them they have four actual stars. With the image on the right, we're using repeated icons to represent repeated instances of a more general class of things and *that* was an unfamiliar graphical convention for our participants. So although these look as though they are in the same format, there are fundamental differences between them that affect comprehension.

An issue that I think contributed to the problems we saw with literal interpretation is that many participants seemed to ignore the titles. There might be that one person in the group who had a college education and pointed out the title to the others and that the recommendation works out to only two servings per day, and then other folks would say, "Oh, I didn't see that." By watching people look at the infographics and following their gaze, the pattern that I observed was that they would look at the image first and then *maybe* consult the title if they felt unsure about what they were looking at. This happened often enough that I started calling it "title neglect" kind of like denominator neglect. I got a really vivid illustration of this phenomenon when I was doing a study comparing infographics to a text-only condition.

On the left is the information provided to the participant. We had a whole conversation about how she couldn't figure out what her BMI was. She knew she wasn't overweight so she guessed her BMI was between 18.5 and 24.9 but she couldn't find her actual BMI value. <click> I said, so what does it say at the very top of the page? And she read me the 2nd line. She had been systematically skipping the first line (even though the font was the same size as the rest of the text) presumably because she assumed it was the title and therefore not important.

What's curious about this is that the phenomenon of title neglect that I observed is directly contradictory to experimental work showing that titles may be what people remember most about a visualization. <click> But I think we can agree that the kinds of visualizations we are showing are pretty different from one another. The ones on the left cannot be interpreted without reading *some* text, whereas mine seem like maybe they'd be self-explanatory. And the research participant populations are pretty different, too. This is a topic that I haven't had a chance to test in a systematic way. It would require an eye-tracking study. How prevalent is title neglect? Who does it? Why? Under what conditions? We don't have answers to those questions yet. What I *had* determined at this point in my research is that when they were inclined to do so, my participants could read accessible text pretty ok. They weren't struggling with the words so much as they were struggling with interpretation, with understanding the implications of what they had read. At this point I'd like to make it clear that designing for people who have limited literacy and for people who cannot read at all are completely different tasks. Designing without words at all is next-level, y'all, ok?

But given that simple text was fine and that I had already settled on number lines as being the best format for a lot of the data types we wanted to convey, I started putting words back in. Here's the front cover of a pamphlet designed to show patients their level of asthma control. My designer, Nicole, and I have done a couple of things here. First, we've essentially put the title on the cover to force people to read it. And we've put a small, redundant title at the top of the inside spread specifically so that they can skip it without missing anything important. Then, we've just given people the punchline. Your asthma's not in control. Your score is 2.5. When I watch people as they read this, I see them start gently nodding their heads. Because all of the cues are pointing to the same conclusion and they have greater confidence that they have arrived at the correct interpretation. You can see the difference that all of these elements make for people. Another innovation that came with this project was successfully (and comprehensibly!) putting two data points on the same number line. We were able to demonstrate that that worked. What participants told me next turned out to be very important. They said, "it's all very well and good to know that my asthma's not in control, but what should I do about it?"

So, on the back of the pamphlet, we left space for the patient or the clinician to write in the highlights of the care plan. We were discussing with them how this section should be labeled, "Your health care provider's recommendations…" and they're like, "No. Plan for you." That's it. Simple. Patients want the bottom line. There have been times when I have asked participants what they want to know about a given health topic. And occasionally, they'll say "everything! I want to know everything!" But I am here to tell you they do not mean it. What they mean is, "I want to know everything that is relevant TO ME and actionable and no more than that." And the reason is that in day-to-day life, adults are problem-based learners. In a formal schooling setting and in cases of particular curiosity (someone's just really into trains, for instance), adults will learn just for the sake of learning. But more often than not, they are motivated to learn in order to solve a real-life problem they have, like how to control their asthma. I studied theatre as an undergrad, including lighting design and a little bit about electrics. Twenty years later I wanted to change a light switch in my apartment. Did I leaf through my old textbook and brush up on basics of electricity? No! I watched three youtube videos and changed the light switch. The point is, when we're designing for lay audiences, we're going to be most successful if we're addressing a topic they care about and a problem they want to solve. Let me tell you about another example from the asthma project.

This is a sample output from a spirometer. It's what's used to test lung function. As you can see, there's a lot happening here. Clinicians require extra training to interpret it because of its complexity. In particular, this chart on the lower left, called a flow volume loop, takes time to really understand because it's got liters on the x-axis and liters per second on the y-axis, which is not at all intuitive. Despite the complexity, my co-investigator, Maureen George, wanted to see if we could extract something useful from here to share with patients. To make this work, we'd have to introduce a little terminology, provide numerical results, visualize anything that might be hard to understand, and help people interpret the metrics. Tall order.

But I gave it a shot. Here's my very first sketch for this design.
<click> Here, we're introducing the terminology for two metrics – FEV1 is how much air you blew out in one second, FEV1 predicted is how much people like you are expected to blow out.
<click> We give the numerical results for those metrics
<click> And then use an icon array to show what that ratio of 71% unobstructed lung volume looks like.
<click> On the bottom right we define the levels of obstruction.
My rationale for using an icon array was that I knew that visualization should be used to ease the most effortful parts of sensemaking. In this case, I identified the part to whole relationship as the part I thought could use the most support. (Spoiler: I was wrong about that.) I didn't have the terminology for it at the time, but I was applying the conceptual metaphor that lungs are a container. The logical consequence of that metaphor is that they can be filled, but in this case, not all the way. We showed this to some focus groups and our participants did not like it. They weren't fans of, really, anything about this sketch but especially the icon array. Because to them,

the dark blocks signified obstruction. Whereas I was trying out the lungs are a container metaphor, they were seeing good is light and bad is dark. We were speaking different languages entirely. They wanted me to reverse the colors, like this.

It still wasn't a winner, though. Far from it. I was able to get my graphic designer on board and together we took a different approach.

Here, our visual focus is on illustrating the three categories to make those easier to understand. But this was still a big WOMP WOMP. Participants were overwhelmed by all the terms and numbers. They really didn't like it.

So, then I thought, what if I present the information in the same way that I would if I were just having a conversation with someone, explaining these ideas and scaffolding the information. Here's an early draft of a visual explanation. After my graphic designer got through with it, it was much improved and looked like this.

Now, we're cooking with gas! It has nearly all of the same information as before, but through layout, size, and color, we've created an information hierarchy. So we're still introducing the technical terminology, but it's deemphasized because it's not the main point. If you read only the large text, you'll get the most important message. And people felt that the belted lungs really reflected their subjective experience of asthma which was way more meaningful to them than some dumb balloons. It turns out that the most effortful part of sensemaking was not the part to whole relationship. People actually feel pretty comfortable with percentages. Rather the hard part was the concept of liters of air. Most of us don't spend time thinking about the volume of air, so that's confusing to begin with, but on top of that, Americans don't know liters. EXCEPT two-liter bottles of soda. So this works because we have converted an unfamiliar metric into familiar terms by leveraging an existing competency – that is, the competency of how big a two-liter bottle of soda is and what it feels like. What's funny is that I had tried out the conceptual metaphor of lungs as a container, abandoned it, and then come right back around to it, just in a different way. But it's not the only conceptual metaphor at play here. See, I was initially reluctant to use this bottle approach because what we're showing is inaccurate: gas expands to fill whatever container you put it in. But I've shown this to hundreds of people and the very first person to complain was actually just this past weekend, because she's a physics professor. She had MAJOR issues with this image. But nobody else ever complains because the depiction is useful. We are willing to apply the conceptual metaphor that air is a liquid because it helps us understand liters of air. I had a participant who had not completed high school tell me that "anyone could understand this." And that was incredibly rewarding. Because getting to this point was not easy.

One of the other things we tried, was to help patients engage in some basic pattern recognition so that they could kind of get the gist of what to look for in a flow volume loop. So we made this explainer for them. We pointed out the features that they should look for. Basically, the smaller and more scooped out the shape, the worse their asthma was.

And it did work. Even when these shapes were neither color-coded nor labeled, our participants could properly classify them. They could get it right. But they didn't care. These blobs aren't meaningful to them. What they care about is their subjective experience of their asthma.

They care about these belted lungs. My participants said, "THIS is what my asthma feels like. You nailed it."

So let's recap some of the lessons learned. We want to use images that are either clearly literal or clearly metaphorical. We want to pay attention to the graphical formats that work in the wild, or that people are experienced with and leverage those. Whenever you can, just tell people with words what you want them to conclude, and if relevant, tell them what they should do about it. Think about the needs of the problem-based learner. Focus on giving them the information that helps them address a problem they care about. Give the imagery that reflects their lived experience. Scaffold the information so that you're giving it in a stepwise manner. Leverage people's existing knowledge and competencies and use metaphors that can help them take in new information.

I hope that you have been able to reflect on how the takeaways I'm giving you might help you in your work, even if your audience is a professional one. The fact is that even professionals with lots of education and domain knowledge and expertise, still have the same basic cognitive limitations as everybody else. For example, your working memory lets you hold maybe 3 to 5 things, tops. And our cognitive capabilities are diminished when we are tired, stressed, sick, etc. The older I get, the more I notice and come to respect my limitations. I'm a very capable person, but probably not nearly as capable as I think I am. If I were to ever write a memoir it would probably be called "Unrealistic Expectations." I'll be happy if you, as the viewer of my talk, remember just 3 of the points that I have made. <click> I highly recommend this book, 100 Things Every Designer Needs to Know About People. It's an easy read and touches on those limitations that we all share as humans, the cognitive biases that we're prone to, etc. Very useful stuff regardless of your audience. Before I wrap up, I want to show you the software that I mentioned before. We created it to tailor that asthma control infographic that I showed you before.

It’s called Browser-based Infographic Tailoring Self-service Interface or BITSI. Here is the interface for generating a single image with BITSI. There's a drop down for selecting the language and a text box for inputting the patient's name. You can select the date for the most recent ACT score. If you want to also show a previous score, you fill out the boxes on the right. If not, you just leave them blank. At the bottom left is a preview of where the scores and arrows will be, and a preview of the text is on the bottom right. If everything looks good, just click Download PDF and it will go to the default download location.

Here's the interface for batch files. You feed it an excel or csv file--there's a template in the BITSI download. And then this first table shows you and rows that have a data error. In this case, June 31st is an invalid date, so we're getting the reason right here in this column. All your good data are summarized next, and then below, not shown, is a table with all the data. When you click generate batch files, BITSI will give you zip folder with pdfs for just the cases that have good data.

I’m sure you’re all very excited to try BITSI for yourself so we have made a demo version available at bitsi.pics – that’s the QR code on the left. I beg you though—do not enter real patient data. It is not HIPAA compliant. It’s just for demonstration. If you want use BITSI for real, if you want the API, you can use the QR code on the right to download it from github.

As reminder, here’s that call for papers I mentioned at the beginning of my talk. Tell all your friends. We want to see lots of great papers.

And with that, I’d like to open it up to questions and discussion!