

On Feb. 6, 2016 an interview took place between Sarah Oppenheimer and William Warren for an article that will be published in the forthcoming issue of aCCeSsions.org. Evan Calder Williams was a respondent, and the interview was co-produced by Laura Herman, Alexis Wilkinson, and Emma James.

The following excerpts have been pulled from an unedited transcript of their conversation.

William Warren: I am a professor of Cognitive, Linguistic, and Psychological Sciences at Brown University, and for about 17 years I've been using virtual reality (VR) to conduct my research. I study the visual control of action of all kinds in the Virtual Environment Navigation Lab, or VENLab. I have been focusing on the visual control of locomotion (how we walk around in a cluttered world) and navigation (how people learn the layout of their environment, and manage to get from one place to another.)

That's where my connection with Sarah really began.

Sarah Oppenheimer: I think that my conversations with Bill arose because I initially had been trying to understand the relationship between locomotion and the material environment—the physical environment in which the body moves—through my work. I had set up a series of experiments where I would solicit people to participate in blind studies, and have them move around in environments.

I would track their movement in space and then, ultimately, compare their movement, as small variables in the space changed. I built elaborate 1:1 scale architectural models and participants were asked to move through customized spaces lined with flexible wall systems. I tracked their movements using multiple cameras, modified the environment in response to locomotion patterns, and ultimately, compared locomotion data resulting from different spatial configurations. The impossibility of accurately modeling an architectural site quickly became clear. This problem of representation was not limited to the simple navigational studies I had been developing. It seemed to be a central question in all experimental models of social behavior.

I became increasingly concerned that by performing the experiment within the framework of an art institution, the form of the test would be taken for the work itself. This troubled me, as it invoked some of the limitations I had encountered in relational aesthetics: the appropriation of social modes of behavior became the form of the work. I wasn't interested in the experiment as form; I was interested in understanding the relationship between architectural surround and spatial navigation.

This led me to research existing studies on the relationship between cognition and navigation. I contacted Bill to find out more. There was a funny congruence here, because what had been developed at VENLab was a very sophisticated virtual model. This model interested me because it seemed to address the ever-present problem implicit in all studies of spatial navigation: how one simulated the space one studied.

WW: I get this e-mail from Sarah back in 2007. I'd been teaching a course called "Perception, Illusion, and the Visual Arts" for many years, as well as other courses on perception and action. I thought, "Wow, this is great! A practicing artist is reaching out to me, and I'd love to talk with her."

Sarah started talking about wormholes, folding space, a taxonomy of holes, which was closely connected to my research. That began a whole series of conversations.

Familiarizing myself with Sarah's work, I saw she experimented with collapsing space, folding space in on itself, and finding connections that would zoom your focus to another space. You could physically walk around these installations, come back to where you started, yet totally fail to understand the space. It seemed like you were seeing one location in space, but when you walked there you'd end up in a completely different location. It was very hard to integrate the experience. This was exactly the sort of thing we were finding in our studies of spatial navigation, and I was coming to the conclusion that it matches the experience people have in daily life. Sarah's work heightens the inconsistencies in everyday experience, revealing the spatial contradictions we don't ordinarily notice.

SO: I think that my conversations with you, Bill, allowed me to develop a language around some of the perceived spatial contradictions that were just beginning to occur in my work.

As I recall, the term "wormhole" was a term that I appropriated from your lexicon. This came about when I first visited the VENLab, and had the extraordinary opportunity to inhabit one of your digital, experiential models. I can describe that model, but maybe it's better if you describe the model of the wormhole. Inhabiting that model was an extraordinary experience for me.

WW: Let me take a step back and say we all have naive intuitions about space, starting with elementary school geometry, where we believe ourselves to be operating in an Euclidian space: it has three dimensions, described by a coordinate system with three axes, which you can mark off in feet and miles. In this Euclidian space, everything is located at a set of coordinates on these three axes. But I've come to believe that our experience of space is radically different from this. In the last twenty-five years, many perception experiments have asked people to make judgments about how far away things look. It turns out that things appear closer than they actually are—perceived space and perceived shape are compressed in depth.

Some of Sarah's work makes this visible. One of my favorite examples is her Mattress Factory installation, where she cut a hole in the floor on the fourth story. It looks like there is a beautiful piece of polished wood sitting on the floor, but when you walk over and look down, it appears that the street is right below the floor. Sarah constructed a wooden chute that runs down through the floor and out a third-story window, so what you're really seeing is the neighborhood outside the building, but it appears to be inside the building, just below you on the third floor. So, Sarah has folded space, turned the building outside-in, put two separate locations in Euclidian space side-by-side. She's collapsed space, revealing that we don't perceive Euclidian space at all.

Something similar happens in navigation, when you ask people to walk around and learn the layout of a space. You may be familiar with the idea of a “cognitive map,” which dates from the 1940s. The idea was that as you move around, you build up a mental representation of space in which each familiar place has a set of coordinates. For example, if you want to go from the grocery store to the post office, you know the coordinates of the grocery store and the coordinates of the post office, so you compute a shortcut that takes you straight from the grocery store to the post office.

Many researchers have conducted experiments in which they ask people to do just that, to take a novel shortcut between two familiar locations. It turns out that people are horrible at this. We found the same thing in VR. We created a virtual hedge maze, made out of virtual shrubbery, containing objects like a fountain, a well, and a bookcase. Participants spent an hour in the maze learning where the objects were. We then asked them to take shortcuts between the places they have learned. For example: walk to the fountain, close your eyes, and take a shortcut straight through the hedges to where you think the well is. They were terrible. Walking in wildly different directions each time, they didn’t really know where the objects were in space.

Then we asked what it would be like to learn an environment that was not a three dimensional Euclidian world. Would you even notice? So, we inserted two “wormholes” in the hedge maze, like that little chute in the floor of the Mattress Factory, so that when you walk through a wormhole you pop out somewhere else. In this wormhole world, when people walked down a certain alley in the maze, they passed through an invisible “portal” and were seamlessly teleported to another location in the maze, and rotated by 90 degrees. The participants spent an hour learning this maze, and we asked them to take shortcuts just like before. People were fine with that. They took what they thought was a shortcut to the well. However, instead of walking to the actual location of the well in the maze, they walked through a wormhole to get there—effectively collapsing space. Rather than learning the Euclidean coordinates of the well, they had learned how to get there, like jumping down Sarah’s chute. They had learned the connections between places more like a network of paths than that of an Euclidian map.

And no one noticed the wormholes. Nobody realized that the wormhole maze was different from the Euclidian maze, or that they were being teleported from place to place. When we asked about this at the end of the experiment, literally one out of a hundred people reported that there was something funny going on in wormhole world.

There are other kinds of geometry besides the Euclidian geometry we learned in elementary school. The opposite extreme is topology, sometimes called “rubber-sheet geometry.” In topology, distances and angles don’t matter—what matters are what places are next to each other, or what places are connected to each other. If you mark some locations on a rubber sheet, stretching it will change all the distances and angles between them, but adjacent locations will remain next to each other. In topology, a network of places is called a “graph.” If you connect places with rubber bands,

stretching them will change the distances and directions, but the links will remain constant. In contrast, changing the links changes the graph. I believe our understanding of space is closer to a cognitive graph than a cognitive map. In some of Sarah's works, where she's thinking about different types of holes and links for instance, I think she is also working with topological ideas rather than traditional Euclidian notions of space.

SO: That is so fascinating to me, because we have never talked about graphs. I am particularly interested in graphs because I recently began developing a new body of work around the notion of an architectural switch. The switch sets up a dynamic exchange in a larger spatial network.

Unlike the Euclidean model, the emphasis in a graph is on the linkage between points. This linkage has no quality. It is not far or near, curved or straight. It is simply connected or disconnected. In this sense, a graph clearly articulates a network—a set of connected (or disconnected) points. A switch alters the link as opposed to changing the proximity. And that is really interesting.

WW: Yes, that is exactly the idea. So I will go one more step. We are finding out that it is not just a link—at least in people's notion of where things are in space. We do have some idea of distance, some idea of how far it is along that link. It does not allow us to locate things in a map, like in the sense that we have no idea where things are in the map. We know how to get from one place to another. We roughly know that it takes longer to go from the grocery store to the post office, than from the grocery store to home. But that's about it. We don't really integrate all of that information into a consistent space or into a map.

So, if you are looking through Sarah's work in the floor of the Mattress Factory, it looks like the ground is right up underneath the floor, but it does not look like it is on the floor. It does in some cases, actually, but you have some sense of distance preserved. It is not completely collapsed so that everything is flat.

SO: Yes. That distance is an interesting variable to play with. Perception ends up doing something much greater than bringing something closer or farther. It actually changes the whole network of links and locations: it shuffles the cognitive map. So that distance becomes very important.

* * * * *

WW: I think of a graph as a mathematical model of our understanding of space, not a mental representation of space that's in your head. Because I'm a Gibsonian, I resist the idea that we have mental representations of reality in our heads. Gibson was reacting against the Cartesian tradition, which treats perception as an internal representation of the external world. This view was born of a deep skepticism that we can know anything about the external world at all. It was, in turn, a

reaction against essentialism, the Platonic tradition that the attributes of things are objectively given and categories are universal.

Gibson is part of a third wave that descends from the American pragmatists William James and John Dewey. Pragmatism argues that reality is neither objective nor subjective, but relational, determined by the ability to act. Our knowledge of reality is determined by the success or failure of our actions. These ideas influenced Gibson a great deal. Gibson believed that we are not locked into subjective mental representations, but neither do we have access to an objective reality. Rather, we directly perceive the “affordances” of the environment we inhabit, which are determined by the relation between material conditions and our action abilities. This is an “ecological” view because each species perceives and acts successfully in its own econiche.

SO: That raises a question about the possibility of a virtual (digital) environment accurately or adequately modeling this kind of ecological environment, one that is greater than the sum of the material envelope supporting and surrounding a physical body.

* * * * *

SO: How accurate is this virtual model, assuming people who are not together experience themselves together?

WW: This is a really complicated question, and I'm not sure we have good answers yet. In VR, one is aware that one is in a virtual world, not the real world, and the question is whether or not this double-awareness affects behavior. We have done some “validation experiments” in which we built a virtual doppelgänger of the lab. We ask participants to do a task in the physical lab, and then put them in the virtual lab and see if they behave the same way. I should point out that in the VENLab, a participant physically walks around in a large 40-by-50-foot room while wearing a head-mounted display, so they have the same embodied experience, and we manipulate their visual experience. So far, the behavior is qualitatively the same. People are a little more reluctant to get close to virtual objects than they are to real objects, but that may evolve as people become more familiar with VR.

I, like many others, get completely sucked into the virtual world, lose track of where I am in the physical room, and start interacting with virtual objects as if they were real. This sense of presence is very convincing. I think it's going to be very easy to accept other people in the same virtual space and interact with them. For example, we're in four different locations at the moment. But imagine that we created a virtual Louvre museum, and each of us has a virtual avatar that is driven by our physical movements and gestures. It will feel like we are all together in the Louvre, and we will interact quite naturally.

WW: There are a number of concerns there. First, I have to say I'm not a big VR booster in terms of the future of human interaction. Second, as you point out, virtual displays have some perceptual limitations, particularly when it comes to distance, scale, surface properties, and the

* * * * *

SO: I'm happy to talk to that but I'd like to deviate from the screen as the subject. I largely think of it as multiple intersecting spatial problems. One problem is how does this interface shape and define the experiential world of the spatial model. How is the model accelerated, flattened, stretched, cropped? Another problem is the model itself. Euclidean models of dimensional space offer a limited range of visualization options. The world can be viewed as a rendered perspectival image, as a ghosted paraline projection, etc. This representational menu limits this expansive world.

Furthermore, the thickness of the screen lies not solely in how it mediates this digital model, but also in how it in turn exists in space, and how it as a form becomes its own portal within a discrete network of linked portals: my screen, your screen, our shared space inside, and so on.

It is also uncanny to consider that these models, and our experience of these models through interface, shape the material environments we inhabit. Our world is designed on screens. The digital model allows for an unprecedented information flow. One party can design the model, one party can manufacture the model (transpose the file into another material), another party can survey from it, build from it, and the resultant material from this chain of social relations is a new architectural landscape. I am concerned this will tend toward an increasing homogeneity in the spaces that we inhabit.

Evan Calder Williams: Yes, it's a revenge of Euclidian space.

SO: Exactly, right.

ECW: Those graphic systems are used both to draw up buildings that will be built and also to model the majority of composite images you see. There is a way in which one gets re-normalized. There's a kind of aesthetic education in the placeability of literally any kind of flake of snow within the graphical world. And I think you're totally right that we come to treat physical spaces as navigable. You know, the old architectural problem of cladding starts to look very different when we come to understand graphic systems that essentially are the cladding of modular surfaces over wire frame, right?

SO: Yes, that's very interesting.

WW: Although with a real surface you can continue to scrutinize and explore. As you do so, more and more gets revealed—the bricks, the surface texture, the micro-texture. But in a model, you hit the scale at which the image is rendered, and there's no more to see, whereas if I'm interacting with a real surface there's always more to see.

SO: Well that's also really interesting Bill because what you had said earlier was that there was a hesitation in approach in your validation model. I was wondering if that would be a result of image resolution. "Resolution" is an odd term because it's not solely about pixel resolution; it also has to do with the problem of focus and focal distance. Perhaps you could elaborate on this.

WW: There are a couple of problems. First, if you put on a head-mounted display (HMD) the screen does kind of melt away. It is a very different experience from looking at a computer monitor or a flat-panel screen—in one sense you are looking through the HMD screen. HMD resolution is getting high enough now that the pixels are barely visible. In addition, because we have two eyes, there are two images on the screen that we view stereoscopically, one for each eye. So our two eyes are not converging on the screen, but on a point behind the screen.

But there is a residual problem: each eye also has to focus on its corresponding image. Thanks to the lenses in most HMDs, your eye has to focus on a virtual image of the screen that is about three feet in front of you. So if I look at an object in near space, my two eyes will converge on a point that is closer than three feet, whereas if I look at an object in far space, my eyes will converge on a point that's farther away than three feet. So you have an inherent mismatch between the focus distance (three feet) and the convergence distance, both of which influence the perceived distance of the object due to stereopsis. This may be one reason that scale and distance are off in VR—things look about 50 percent closer than they should.

...

WW: Well, in a virtual environment we have a ground plane, right? I can look down in my HMD, and there's the ground plane, but I don't see my feet on the ground plane, so it is disembodied. My suspicion is that when we can see our virtual feet on the ground plane and we can solve this mismatch between convergence and focus of the eyes, the scale problem will get resolved.

The theme of this issue of aCCeSsions is scale. The ground plane establishes the body scale for humans and terrestrial animals. Without that ground plane and body scale, you have no idea how big or how far away anything is. Size is scaled to our bodies. Body scale is given by our eye level. If I stand on the ground my eye is just over five feet above the ground, and that sets the scale for all the other objects on the ground plane. If I look straight ahead and I see a telephone pole, my line of sight cuts the telephone pole at my eye height, and I see the telephone pole as three times my eye height (because it extends three times above my eye level). And so, as was understood when perspective was invented in the fourteenth and fifteenth centuries, the eye level sets the scale for our perception of the size and distance of environmental objects.

SO: That problem also pertains to the possibility of sight in institutional exhibition spaces. The scale of the exhibition space, especially in museum architecture, functions as a lucrative urban planning decision. The scale of these spaces becomes increasingly divorced from the level of the eye such that, as you enter the museum space, the scale of vision does not match the location. This is further compounded by how these spaces circulate. It is not just that we circulate through the exhibition spaces—representations of the exhibition spaces circulate through a secondary level of distribution, which is film and photography.

Documentation of these spaces further compounds this condition: in photography, eye level is actually not attached to a body. This gives the photograph the capacity to rescale the architectural surround. There evolves a demand for works to be scaled to the photograph, in a way that divorces them from embodied experience. That then raises further questions about the difference between eye level in a mediated environment versus in the embodied, corporeal world.

* * * * *