Visualizing opportunities

interview with Chris Harrison
1.14.10

As designers we are primarily familiar with your mapping and visualization projects. However, your work goes deeper than that and seems to engage several interdisciplinary studies. Can you tell us about your work with the interface between computers and their users?

Sure. My formal training is as a computer scientist. I went to New York University, got my bachelors and masters degrees in computer science . So, I am sort of a traditional programmer on one end - and then, after working in the industry, I came back to work on my Ph.D. in a fairly new field called Human-Computer Interaction (HCI), which is often described as the intersection of computer science, design, and cognitive science or behavioral science. It's a weird intersection between quite distinct disciplines, with the common goal of enhancing the user experience.

You can't just create technology in a vacuum. Computers have followed Moore's Law, but humans haven't. There's no point in making computers faster, and faster without addressing utility. People haven't gotten any faster since the advent of computers. Our performance has been constrained. And so, you have to take all these distinct disciplines, like behavioral or cognitive science, design, and computer science, and work with them in concert to create interactions and experiences that are more fluid, more powerful, more capable, and so on.

This is what I spend my day thinking about. In particular, my interfaces work, my core research, which is what I publish on and go to various international conferences to present, is looking at two distinct problems. The first one is interacting with small devices in big ways. Computers have been able to miniaturize fantastically, mostly because of advances in electronics. But, you can't simply make devices really small. The iPhone is a great example of this. The iPhone is the size it is, not because Apple couldn't make it smaller - they could pack the iPhone into a wristwatch, but how would you effectively read your emails and do pinching gestures on something the size of a wristwatch? What I look at is how can we continue to miniaturize devices, but make their interaction as fluid, and capable, and expressive, as full-sized systems, like your desktop.

The end result for this might be devices the size of a wristwatch, but you interact with them at a much bigger scale. You are not constrained to the physicality of the device; instead, you can gesture with your hands, or maybe you projection onto the environment. I just did a project this past summer, called Skinput, where you just use your hands and your arms, and really whatever body part you like. You can actually "click" on your body with a finger, and the system, which could be hidden anywhere on the body, listens to the signals and actually knows if you taped your palm or your middle finger or your pinky. Imagine if you had an iPod strapped onto your arm and you wanted to go to the next song or change the volume. This could be done without having to reach over and pressing a very small button. You could actually gesture right on the palm of your hand. There really is no physical interface at all. You are blurring where the device begins and where it ends.

So it's no longer about being faster, it's more about the experience. Where can we take this?

Well, that is the million-dollar question. So, you remember the Motorola Razor? That's a six-year-old technology, and we really haven't made phones that are much thinner than that. After a certain amount of reduction, it becomes impractical to have any smaller of a device. The device form factors that we have today, like the MacBook Air, or netbooks, or iPhones, have kind of plateaued in size. We are not necessarily going to get them any smaller unless we come up with new ways to interact with them. And it's not clear whether you would really want to call your boyfriend on your wristwatch phone by writing numbers in the air, you know? That's not quite as intuitive as pressing buttons on a screen.

I think the open questions are either how do we come up with new interactions so that we can continue to miniaturize, or how do we make the device size we have today more capable. The iPhone is a paradigm of the latter. It didn't really have any technology in it that was spectacular. The real success story is how its creators thought: 'How do we create a really nice user experience and package it in a phone that people want?' The iPhone is actually quite bulky compared to some other phones out there. So it's clearly winning with its interaction.

The reason people aren't jumping ship to other platforms, like Google's Android, is because the user experience is not quite there yet. The iPhone really has set the gold standard. I have a Nokia phone that I just hate. It's so hard to use! I don't know 90% of its functionality because it's just so convoluted to access. To really make these devices better, the human interface has to be improved. They don't need to be faster or have a higher resolution camera. They need to improve the user experience. And I think that's obvious to everyone. The iPhone was a wake up call to the entire industry, saying 'OK, we made these phones with a million features, and people didn't use half of them, and didn't even fully understand how to use the ones that they had access to. The iPhone didn't even ship with the ability to MMS, yet it made huge inroads.

So, maybe we can talk about how you create this user experience, technically speaking. How do you use all this information that you compile and show people how to use a device intuitively? And what are some of the challenges you encounter in designing these user interface systems?

As I said in the beginning, there are two roads. On one hand, you can keep devices the same size, and make the user experience on that platform richer, or you can come up with new ways to interact with devices, and let you miniaturize further. I'm currently researching the latter. What I try to do is find funky sensing solutions. I add extra sensors to a device so that I can explore new abilities. An example is a project called Scratch Input. Basically we experimented with a microphone that could potentially be placed in

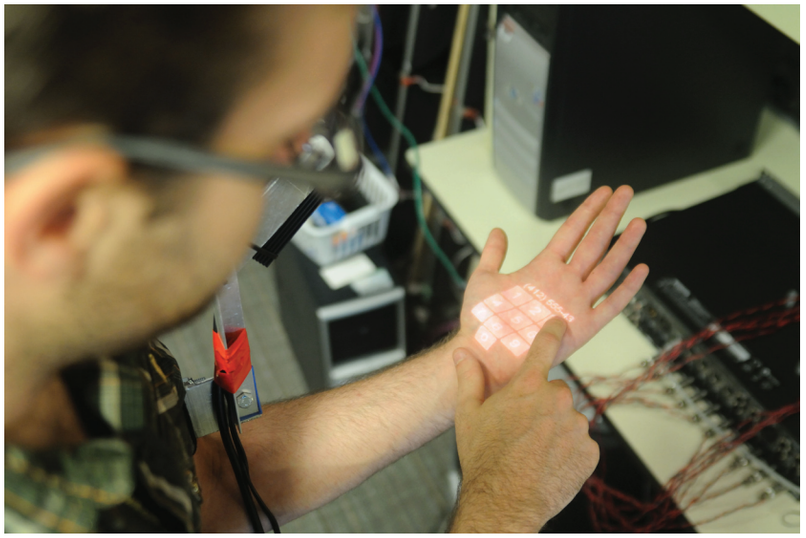

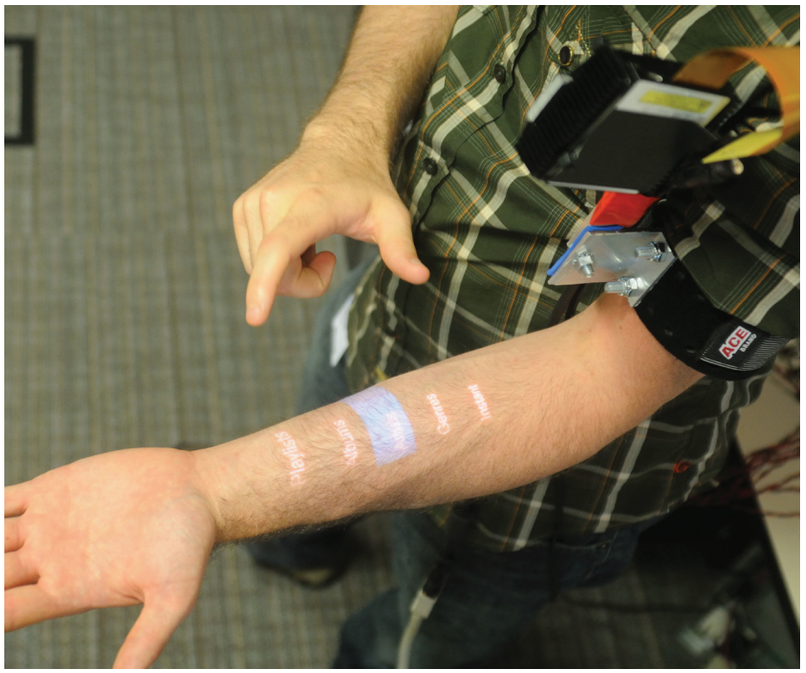

Freed from the physicality of a device, Skinput appropriates the body as an input surface.

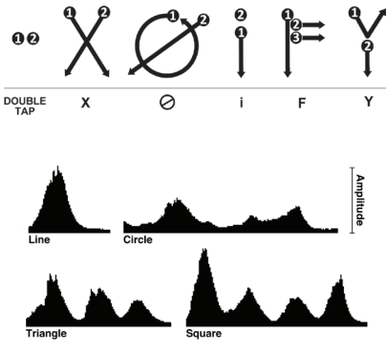

Scratch Input uses a small, inexpensive sensor to track gestures made on a surface, and communicate them to the device.

the bottom of your laptop or on the backside of your phone. When you put it down on the table, the microphone is pressed against the table's surface by the weight of the device. Now, instead of having to press the tiny buttons on the phone, you can actually just drag your fingernail on whatever surface the device happens to be resting on. The device listens to the acoustic information in the surface.

Essentially, you can gesture to your phone on almost anything. Imagine when you are at work, you have your iPod, there is a special microphone in the back of the iPod, and you plopped it down on the corner of your desk. You are listening to your music and want to go to the next song. You could reach over and press the next song button, or perhaps draw an arrow right on the surface of your desk. The device can hear that motion and go to the next song. What's really cool about this is you can have a really small device essentially steal input surfaces from its environment wherever you go. So, I can put it down in the airport, or in a Starbucks, or at my office, and I get a huge input surface in the form of a table. Of course, the device stays the same size. Amazingly, the scratch input sensors costs about a dollar – so this isn't expensive technology.

I attached one of these sensors to a wall in my house, and built a really simple music player. The idea was that I could walk up to my wall at the end of the day and double knock on my wall to play music. I could make a 'V' for volume mode or 'S' for seek mode. I could make a circling gesture, like an iPod, and change the volume. I'm just writing with my finger, like a passive touch screen on the wall of my house, which is just a wood frame house. I can take my house, which was never designed for input, and all of a sudden infuse it with computational ability with a one-dollar sensor. Pretty neat.

What do you think the lag time is for technology like this to reach the market?

Probably like 15 years. Again, jumping back to the iPhone, the iPhone did nothing new - multi-touch is something like 30 years old. There were people demo'ing things that the iPhone was doing 15 years ago. People who are in the know say, 'Oh wow the iPhone, I remember this paper from ACM CHI 1985,' or something like that. In academia, research labs tend to be really ahead of the curve. Of course, when they are doing it, they're doing it on an iPhone the size of a small refrigerator. That's how one mocks up the future.

It's not commercially viable for years, and then someone actually invests all that effort, demonstrating that you can do multi-touch and those really cool pinching gestures and so on. All of that was done a long time ago, but creating a polished product that isn't going to make consumers rip their hair out can take years, even decades.

I believe Apple was working on the iPhone behind the scenes for more than five years before it came out. It took competitors… I mean the iPhone came out in what? 2007, 2006? Only now are we starting to see companies with real competitors to the iPhone. So you're talking about at least 2 to 4 years of lag, even if you get right on it with a lot of resources. And academia is often way ahead of manufacturers, but the prototypes are crude. I can do 10 projects a year that are really cool, but to actually do the last 10% of the work to bring them to market is 90% of the effort.

Going back to what you said about leaving this technology open to different users, have you ever been surprised by the way people adapt these technologies?

Occasionally I'll get goofy emails with people using them in weird ways. Like for Scratch Input, I had an Indian student e-mail me who was trying to make a really low cost ATM in places that did not necessarily have enough power to run a touch screen monitor. Also, because touch screens would get damaged or stolen, she thought of having a printed interface - perhaps an embossed metal surface - and then use Scratch Input as the means of interaction. Again, it uses a cheap sensor to recognize buttons and could do authentications by perhaps writing your signature. It's a clever, low cost, ad hoc kind of ATM machine. And it might just work. As academics, it's takes so long to get things out into the world. So when people pick it up and start utilizing it, it's really rewarding.

Academia unfortunately is its own little world. Academic texts aren't super accessible to regular people. This is why I'm a big proponent of publishing YouTube videos concurrent with academic papers. It's a much more accessible format that people can share more rapidly. A former student here, Johnny Lee wrote several Wiimote hacks – especially head tracking. He had the most popular video on YouTube of all time for something like six weeks – an impressive feat! That demonstrated to me the value of putting academic projects out there to get people excited about them. As technologists we want to see people getting excited about technology. That's what makes us excited. We're not just inventing technology for the sake of technology. That happens so much already, and a lot of devices are terrible.

'Wow, why are we using mice? Why aren't we putting our fingers all over the displays

To get people excited about technology, want to invest in technology, love technology, and hopefully even be educated in technology, maybe even become HCI researchers, is what makes me sleep happily at night.

Do you think getting people excited about it or putting ideas on YouTube would reduce the lag time with the technology transfer to the public?

That's an interesting idea. I don't think that YouTube has been around long enough to really quantify how much it speeds up the process, but the great thing is that when people get excited about a project online, it does put pressure on companies to implement those features. If companies think 'hey, if this grad student at CMU can build this really cool system, why can't we include that in our next release?' This sort of does happen. Johnny Lee, who did all the Wiimote stuff, ended up getting hired by Microsoft and is working with the Xbox group.

I think YouTube will have a measurable impact on where companies take their products and it's a great way for them to get feedback on their products. The Apple iPhone was not the first multi touch device. The first popularized demo was Jeff Han's multi-touch work, which was presented at TED. That kind of kicked off this multi-touch firestorm, even when Apple had been working on multi-touch internally. They just hadn't said anything publicly. After the TED talk, people were like, 'Wow, why are we using mice? Why aren't we putting our fingers all over the displays and pinching and painting with multiple fingers and having ten people use a screen?' It seems so obvious and once people understood this concept, they thought 'why don't we have these fluid interactions, and why don't we have them for all of our computing systems?' Apple happened to have a product already in the pipeline and was able to capitalize on that momentum.

You mentioned earlier that your work is a combination between computer science, design, and behavioral science. And you yourself work in a variety of media and have a considerable amount of experience in different fields. You've worked with Microsoft Research, IBM, and have even been on expeditions to Turkey. How would you say that this sort of diversity in your interests has affected your work?

A lot of people in my university go to school for a pure computer science education. Then you have the bunch of misfits that are doing human-computer interaction. It just so happens a fair amount of us are computer scientists, but we don't spend our time really contemplating computer science problems. We say, 'how can we use computer science as a resource and skill?' I think that kind of mentality attracts a much more diverse group of people - people who are more anthropocentric in their outlook on life. A lot of people in the program have traveled extensively or have lived overseas for decent periods of time. I think what you've touched on is important. I can't totally generalize, but a lot of the people in this department are very experience oriented, which makes sense, because technology is such a central part of human society. It's all about the experience. The function of the device doesn't matter. It can have all the power in the world and have a million functions, but it doesn't matter if the experience is broken

So, I think HCI people approach problems in a slightly different way. Often they come from art, computer science, design backgrounds - one guy was even a carpet seller in Marrakesh. It's a really eclectic group. People coming into this weird field tend to have hobbies like ceramics, and paragliding, and unicycling, and juggling. It's a funky collection of people who approach problems in funky ways, and that's what you need. It's not an engineering problem to make computers better. We're not going to engineer our way out of this problem. We need to turn to interaction. Computers are fast and humans are fast; it's the human-computer interaction that is lacking.

What would you say from outside your normal work, has most influenced your work?

That is a very good question! What I always find interesting, when I travel, is that a lot of the world doesn't rely on technology. It's not like you go to a farmer in Turkey and say, 'You really need to get a web server in your barn.' That doesn't enhance their capability. Technology for its own sake is not good. In the US we have a very gadget-driven culture, which can make it hard to see how technology will impact other people's lives. Clearly, it's a driving force, and I'm sure it will continue to be, but the US is not the rest of the world. We're not going to all have little robotic puppies playing with our kids. I think having a broader view of what technology means to different populations, and how they view technologies as tools and social platforms, is a really interesting question with really interesting ramifications. You certainly have to consider that when you're producing something. Traveling is a key way of opening your eyes to how people use technology in other countries.

On the other side of things are personal experiences. I did dance as a kid and ceramics through high school and college. Ceramics is this visceral and tangible experience. Your hands re directly involved in shaping. Computers, at least how they are now, are far less direct. Multi-touch is starting to move in this direction, but with a mouse on a laptop you're decoupled. You're not pointing directly at the thing of interest. Instead, you have this device that's marshaling your input into the computer. You're constrained by the affordances of the track pad. You can slide your finger in two directions, but when you think about how many ways can your finger move? Not only can it move in three dimensions, but you can twist your finger, you can bend your finger, there's different postures, you can use it up-side down, you can apply different levels of force. A human finger has an incredible amount of affordances that are being redacted to this 2-D plane that your computer

and pinching and painting with multiple fingers and having ten people use a screen?'

understands. Ceramics is a hands on experience. There aren't many gadgets these days that can be so directly manipulated. How can we bring that kind of interaction to computers? Do we want to? Is it an advantage?

Actually, if we look at projects like Microsoft Surface, they're starting to approach this vision of direct manipulation. I put my finger on a photo and I can stretch it like a rubber band, and I can drag it and pop it in my phone, which is lying on the table. That's obviously a much tighter coupling with the physical world than touching a track pad. A track pad works great for word processing and clicking on hyperlinks, but would you ever hope to shape a pot or arrange flowers on your laptop? That's a task that we have no concept of doing.

Interactive surfaces are opening up all sorts of doors. Huge communities of people are painting on their iPhone while sitting at a bus stop and posting these pretty amazing, first-rate illustrations. I mean that's incredible! Professionals do this in Illustrator with a Wacom tablet and a $3000 machine and now people are doodling on their iPhone and producing fantastic images. And the iPhone's not even a particularly good platform for painting.

We're never going to word process on our table. The keyboard and the laptop are really good at typing. We're not going to see people all of a sudden writing essays on their iPhone, but they'll be doing different things. I've seen some funny things recently. A friend was sitting at her desktop computer and wanted to check Facebook. To do this, she pulled out her phone and used the Facebook application. Why? Because it's easier than using the desktop computer she happened to be sitting right in front of. We're seeing that devices are specializing and getting better at particular tasks.

Along that note, as designers we always have to sift through a massive amount of information. I think this is why we like the visualizations you have on your website. It's a method to understand the way we're living as individuals and as a society – and as architecture students that is something with which we constantly grapple. How do you see this graphic representation taking us further, or is it simply interesting correlations?

The visualizations, which are a totally different dimension of my work, act more like a hobby. My computer science background a skill, is much like architecture. It gives you some primitives for building things and understanding the systems of buildings, and you use those skills to produce design artifacts. There are obviously a lot of opinions on how that process works, but that's how I see it. So, having access to these fantastically big data sets, I had a revelation in how I wanted to do information design.

A great example of this are my Bible visualizations. It's a really simple premise. I had this enormous table of cross-references – so, for example, I could see that chapter 5, verse 2 and chapter 10 verse 12 were connected. In its aggregate form, it was kind of useless. So, I said what I'm going to do is to give each data point a really simple visual primitive; in this case an arc. I basically made one function, that must have only been about ten lines of code, and it said: given two chapter locations, draw an arc between the two. When you repeat that tens of thousands of times, even though you've used a very simple visual metaphor, you have complexity because the data has complexity. All of a sudden you weave this tapestry that is really nothing more than a lot of arcs. I like to think that I am not really the designer in all this, but rather the data is its own designer. I'm just giving a voice to the data, so that it can draw itself. I quite like this romanticized view of it. There is a bit of aesthetic tinkering to play with, but the complexity and the beauty in a lot of the work that I do has nothing to do with me. The beauty is inherent in the data, and that's really exciting.

In the case of the Bible visualization, the process only took about an afternoon to get the initial version working. I polished it a little bit more the following day to get basically what you saw. Essentially the whole process took not even twenty-four hours. And if you look closely enough, there's actually some really interesting patterns that you can see visually between the Old Testament and the New Testament.

There are a lot of people doing really amazing data visualization and what they label as infosthetics. I don't do it nearly as rigorously or as frequently as other people, and so can't comment on it so much, but it is an interest of mine.

What is it about this process that interests you? Is it that you don't know what the data is going to look like? Is it the act of creating or just exploring these different sets of information?

I guess it's all of those. I haven't gone out too many times to find the data. Often, someone will e-mail me saying, 'Hey I just compiled all of these cross references in the Bible,' or "Hey, I just got a hold of twelve years of internet searches," or something and don't know what to do with it. I say 'Okay, let me think about how I want to visualize it.' Part of it is just the fortuitous nature of stumbling onto an interesting data set, and part of it is deciding how to visualize it. The Bible has a more linear progression, so arcs and a straight line from beginning to end make sense. Other data sets may have a lot of different ways you could approach it. So, there's a challenge component to it. And the third one is the surprise.

I think the most complicated visualization I've built is only a few hundred lines of code. Whereas something like Microsoft Office is millions of lines. My programs are primitive by most peoples' standards. It's just nice to say I took a gigantic text document - full of numbers - and out popped this rather remarkable thing.

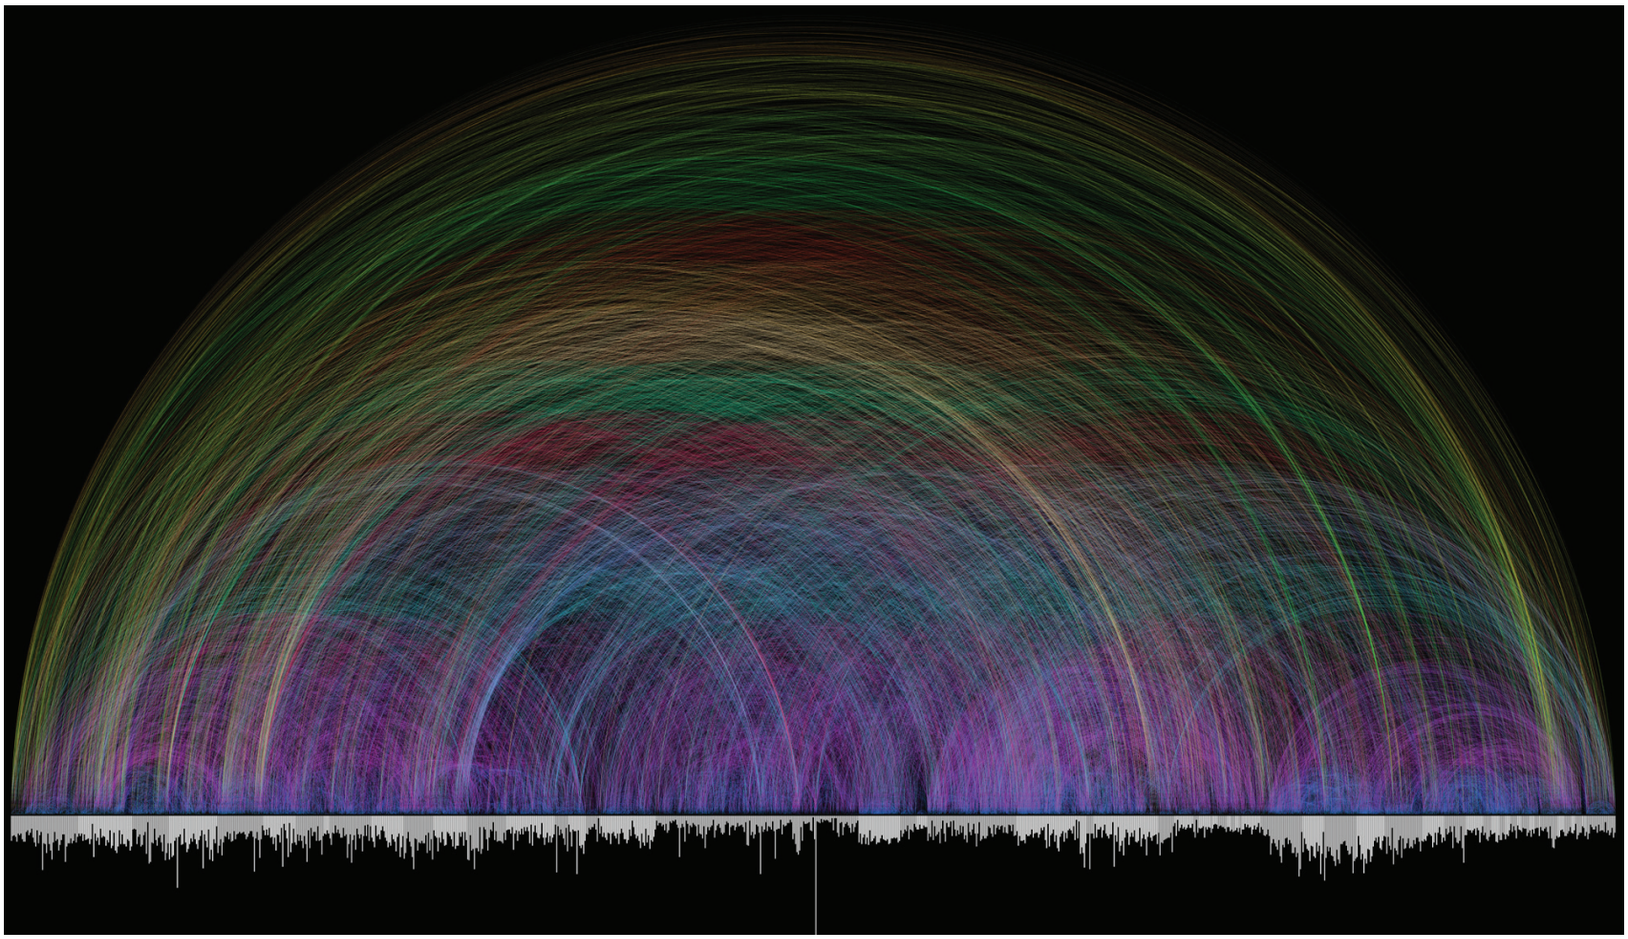

'Bible Cross-References'. Visualizing connectivity in the Bible through textual cross-reference. There are 63,779 cross references found in the Bible, and each of these is depicted by a single arc - the color corresponds to the distance between the two chapters, creating a rainbow-like effect.

Have you been surprised by any of the things that you find at the resolution of these visualizations? Not just by what it looks like, but coming across any trends or if seeing the data in this way reveals any new information to you?

Take, for example, the Bible visualization. It has a popular appeal. People e-mail me quite often saying, "Oh my word! This is really amazing," "What's this chapter?" and "Why is that connected to that?" People more religious than I have drawn great meaning from this visualization. In fact, someone has produced a derivative version of this called Contradictions in the Bible. It looks exactly the same as mine, except a slightly different color scheme, but it links all the pieces of the text that contradict one another.

Other visualizations, like my Internet map, illustrated the incredible disparity between connectivity across the globe. When I rendered it the first time, I was incredibly surprised at the disparity. There is also some work I did later with a Google data set that I got hold of. These are my word association and word spectrum visualization. These are basically a weighting of how words are used in different ways. One is fan shaped, where two words are pitted against each other, like hot and cold, or past and future, and you get to see what words are associated with different things. It's interesting to see how language is being used and coupled with different meanings. What is coupled with woman or man, and

why are these words associated with women but not with men? There are obviously some very comical things in there as well.

You can also do things between nations. I could pit the words "China" against "America." What are people saying about China? What do they associate in the English language with China, but not with America? What do they say about Steve Jobs but not Bill Gates? I thought this was pretty interesting, so I produced a dozen or more of this visualization with different word combinations . I wouldn't say there are any ground breaking revelations, but it has been an intriguing experience.

One thing that we find interesting is that the internet is continually growing and changing, so if you were to do another visualization with updated data, how do you expect it would change, and do you have a particular one that you would be interested in recreating?

How things change over time is really interesting. I would like to redo a lot of my visualizations, especially some of the older ones to see this effect. People are requesting that I do this, so we can see how the dynamics changed over the past two years (since 2007). The temporal dimension is really interesting, and the new generation that is going onto the Internet is very different from the people using it now.

A really tricky part of data visualization is getting a hold of the data. Data is everywhere. Making it cohesive is actually much harder than it may seem. One project, for example, is called "Search Clock" (something I worked on for AT&T). Those guys just happened to have data sets that they collected themselves. For example, they took a snapshot of what was searched every ten minutes for four years on the search engine Magellan. So, I did this temporal visualization with a clock metaphor. Some really interesting insights came from this on what and how people were searching from 1997 to 2001. You could actually see how the searches changed over the course of the day, as well as how they morphed over a four year period. I drew a lot of personal conclusions from this data. People were searching for things like web chat and midi files in '97, and by 2000, they were searching for free software, porn, and information. So the tastes in the web population have changed dramatically.

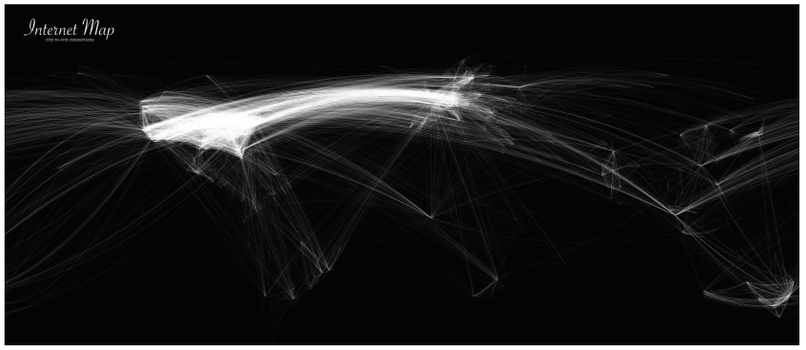

'World City-to-City Connections.' Mapping the connectivity of the Internet around the world.

On the potential of the Internet and the use of this graphic mapping strategy, do you think now that we have these visualizations we can start to understand just what the potential of the Internet is in this collective community? Can you elaborate on the potential of these graphics and how we can use them?

Yeah. This is a little outside of my field, but if you look at a site like Wikipedia.org, the masses have an incredible power to produce incredibly and beautiful things. There is another funky project that I set up with a friend called "Open Mosaic" (www.openmosaic.org). It's basically a mosaic that you can visit and fiddle with the colors of tiles. People go onto this website and make fantastic images. Partially, this is happening in a collaborative way.

It's somewhat like hundreds of people walking up to a blank wall with a fingerprint of paint. You might imagine it would look like static or something random, but humans don't work that way. We are creatures that love to organize and fiddle; that is just in our blood. So, if you told everyone to dip their fingers into any number of colored paint buckets and walk up to a wall and press their finger against it, people would probably start drawing suns and trees, and smiley faces and hearts automatically. That is just how we operate. We don't enjoy randomness, and we don't enjoy entropy.

We can also look at anonymous situations like Facebook, with a population in excess of 300 million. That makes it the fourth most populous country on the planet, which is just unbelievable! We are talking about Facebook, the company, governing a population and defining rules for that population that's essentially as large as America, but really cover the globe. That's not something to be taken lightly. If that community assembled, what kind of power can they wield across the globe? If all of a sudden you have 300 million Facebook users vote for this guy, or vote for that policy, or want free speech here… that's like everyone in the United States saying, "We want a third legislative branch" or something huge. That's a serious amount of power that this population wields.

So, as knowledge continues to change and grow, do you think it's important that we have these sorts of historians of the Internet as a way to see how things have changed?

That's also a really huge question in the science community. If you look back at Newton and Darwin, for example, some of the best records we have of their thoughts are wonderful written letters to their colleagues. That is a huge wealth of information. Unfortunately, with e-mail, most dialogue is just gone. And similar is true with the Internet.

There are efforts underway to archive the web. There is the Way Back Machine (web.archive.org), which has been saving web pages for some time. They are doing a great job, but there is so much out there, it's never going to be 100%.

I think e-mail, which is now replacing all of our letter writing, is actually the saddest thing. There are probably so many ground-breaking discoveries - artists and scientists making their breakthroughs and dialogues shaping the course of the world - all happening through e-mail. We'll never have an archive of this.

Whereas letters from Napoleon and Washington are archived for us to read and contemplate. Letters are essentially lost from our generation forward. It will be interesting to see what the historical ramifications are for that.

To conclude, as a tool of reflection, is there a way you can see your other projects related to these visualizations, and can we use these things to see what sort of influence our work has had on the world?

I think it's certainly an interesting problem. I am glad it is so large and amorphous now, but it's not clear how one measures impact. There are the crude metrics such as YouTube viewings, but there are only so many that go on YouTube. It's really hard to see what is significantly impacted any more, and to visualize it is an even harder. I don't know if I have a good answer, but certainly when you're talking about millions of anything, visualization is an essential component.

If we can move numbers and words and relationships into visual analogs, then we can take advantage of the incredible parallel processing that our eyes and visual system afford. We can see and process a huge amount of data in a very intuitive and powerful way - and that's the key. Align data with how our visual perception works. In particular, with colors and orientation, shadows and shapes. If we can leverage these visual attributes for data visualization, there might be some really powerful things we uncover.

I think this embodies a whole movement that could explode. Data visualization is big. Within the last decade, humans have developed the computational power to create really large data visualizations The biggest influx of people contributing to this field are still amateurs, which is kind of one of the tag lines for your issue. If you go to web sites like infosthetics.com or visualcomplexity.com and look at the communities centered on information visualization, like 90% of them are tinkerers and hobbyists. They may be amateurs, but they have been given incredible power. You now have people producing really complex things in their spare time. And computers have enabled that to happen.