A Conversation with CHCCS 2021 Achievement Award Winner Tamara Munzner

Tamara Munzner
University of British Columbia

Abstract

The 2021 CHCCS Achievement Award from the Canadian Human-Computer Communications Society is presented to Professor Tamara Munzner (UBC) for her numerous high-impact contributions to the field of information visualization research. Her research interests include the development, evaluation, and characterization of visualization systems and techniques from both user-driven and technique-driven perspectives. CHCCS invites a publication by the award winner to be included in the proceedings, and this year we continue the tradition of an interview format rather than a formal paper. This permits a casual discussion of the research areas, insights, and contributions of the award winner. What follows is an edited transcript of a conversation between Tamara Munzner and Matthew Brehmer (Senior Research Staff at Tableau and former PhD student) that took place on March 8th, 2021.

The Interview

Matthew: Hi Tamara, congratulations on being recognized as the recipient of the 2021 CHCCS Achievement Award!

Tamara: Thank you. I'm honored and delighted!

Matthew: I'll start off by stating that you've had a really fascinating career so far, and while many readers might be familiar with some of your accomplishments in visualization research, I imagine that few know about your origin story. How did you first become interested in visualization?

Tamara: The origin story is interesting because it was pattern matching. I was looking for a summer job. I was working at a supercomputer company called ETA systems, a spin-off of Control Data, during my first three years as an undergrad. When ETA closed down, somebody pattern-matched my résumé to the [University of Minnesota's] Geometry Supercomputer Project because they saw the word *supercomputers*.

This pack of mathematicians had requisitioned supercomputer time to draw tilings in the hyperbolic plane. Their algorithms were exponential, involving huge amounts of competition, so they had access to Crays. They came up with some lovely theory, a mash-up of group theory and finite state automata, which they called automatic group theory. They figured out how to do these drawings in polynomial time, so they didn't actually need so many Cray cycles. They used graphics workstations as front-ends to access the supercomputer through, and they realized that there's a lot of interesting issues in drawing these pictures, so they morphed into a visualization research group. The Geometry Supercomputer Project later became the Geometry Center, which was the short name for the National Science and Technology Center for the Computation and Visualization of Geometric Structures.

I fell in love with visualization. My first summer job at the Geometry Center was to take an interactive 3D viewer called *MinneView* (which was substantially written by Pat Hanrahan) and to add RenderMan functionality, so that you could export high-resolution images. I ended up learning graphics by reading Pat's source code and the *RenderMan Companion*. This was quite an exciting odyssey of discovery, and I decided *"Ooh, I like this,"* so I went back and worked there for four years after graduating.

At the Geometry Center, we looked at non-Euclidean geometries and figured out out to project from 4D and 5D down into 3D. My first job title was *apprentice*, a great title for people who didn't know what they wanted to do with themselves. I didn't know if I wanted to go into industry or academia. I thought that I might like academia and that's part of why I took that job: I thought *"I could go write compilers at Sun, or I could go do visualization with this bunch of mathematicians; ... I think I'll pick the latter."*

I got very enchanted with visualization, doing software development for interactive 3D and 4D visualization on a system called GeomView [26]. We made movies to bring mathematics and topology to a general audience, where we used computer graphics and visualization along with non-technical language to get people to think about ideas that usually you'd only see in a grad level math class; the ideas were accessible if you didn't get sidetracked by all the notation. We made videos about turning spheres inside out without poking holes or creasing them (*Outside In* [14]), or about life in a space that's finite that has no boundary (*The Shape of Space* [16]).

I loved my work at the Geometry Center, and there were many good things about it, but I realized that I was doing a mix of research, system administration, technical support, and video production. If I wanted to do research anywhere else, I probably needed to get a PhD. I read a whole bunch of SIGGRAPH papers to figure out whose taste in research I liked, and I ended up coming back to Pat Hanrahan, who at the time was at Princeton. I discovered that he was about to jump ship to Stanford, where I had done my undergrad, so I ended up back at Stanford to do a PhD with Pat.

Unusually enough, while I did a PhD with Pat, we never co-authored a paper, which is extremely rare, but I learned a huge amount from him about the process of research and how to write papers. He was an amazing mentor to me.

I originally thought that visualization was a useful folly, and that I would join the grown-ups and focus on computer graphics. But then I did one little side project in visualization, ... and then I did another side project, one that actually turned out to be quite substantial and produced multiple papers, ... and I ended up getting a consulting job at SGI to productize it [19]. By the time I did my third side project, a system for computational linguistics during a summer internship at Microsoft Research [24], Pat and I decided that it was time to admit that the visualization side projects were in fact the topic of a thesis, so I stopped calling them side projects.

Matthew: Your engagement with industry research has continued throughout your career. You mentioned SGI and Microsoft, and before joining UBC, you worked at the Compaq System Research

Graphics Interface Conference 2021
27-28 May

Center. I'm curious to hear your take on what makes for a good industrial-academic research collaboration.

Tamara: Thinking about my two years at the Compaq SRC (formerly DEC SRC), what I really liked about that environment is the breadth that people would have there. They wanted to do interesting things that had never been done and to have an impact in the research world, but ideally also some impact for the company.

It's a very different environment when you're inside a lab compared to when you're a faculty member collaborating with industry. At a research lab, when the times are good, it can be more fun because you have very low overhead. If you wanted to just do research, you could spend 95% of your time on it. But suddenly, the overhead could could ramp up and instead of spending 5% of your time on overhead, it could be 50% or 90%. Anybody above you could kill a project by saying it doesn't align with the company's mission. The inverse is true as a faculty member, where if you can find anyone in the world who's willing to fund a project, then you can do it. There's an inverse proportion of autonomy as a faculty member, where you have a constant (at least 60%) overhead and that will never go away. On the other hand, they're never going to decide to cut the computer science department out of a university. Whereas in industry, they could decide to cut the lab, or at least cut the research area you're in, such as visualization.

The projects that I like the most are these grassroots projects where there's me on one side and there's one researcher or someone who's research-aware as the main contact on the company side. Ideally, we co-supervise a grad student who's working on the project; we find a way for everything to align. The project has to have research value so that the student and I are able to publish papers; the industry person may well be a co-author. We want the company to care. Ideally, if they're funding the project, they should be convinced. It's an example of *problem-driven research*, where you are trying to solve a real world problem: to go all the way from this domain-specific problem and thinking about it abstractly to designing a system and implementing it. What's nice about working with a company is they're likely to have a really clear understanding of what their problems are. Sometimes, they are the only place you can get certain kinds of data. We've done a interesting projects with quite proprietary data that you could never get access to if you didn't collaborate with the company. Especially for projects that scale, you're just not going be able to gather that much data as an independent researcher. I've had such collaborations with AT&T [17] and Tableau [25] and I've really benefited from that.

Matthew: Looking back on your problem-driven research projects, whether or not they had industry partners, which would you consider to be your most significant contribution or favorite result?

Tamara: That's like asking someone who their favorite child is. But I could speak to ones that were the most surprising, or that I learned the most from.

When I think about problem-driven research, I have a soft spot for the biology and genomics projects. For one thing, the domains are just so cool. The biological term is *model organism*, but I always think of them as *pet beasties*. I've gotten to interact with people who have different kinds of pet beasties: strawberry guy, jumping spider guy, fungus lady, chicken fish guy (who studied both chicken *and* fish), and this person who studied extinct rhinoceri. It's incredibly interesting to get sucked into the world of these biologists, to figure out how to build a tool that could actually help them. What's great about biology is that the concerns of fruit fly lady turn out to be very relevant to the concerns of people studying large vertebrates, because there's so much in common with the genomics of them. Biologists are interesting because they often have big data, clear questions, quite a bit of funding (by computer science standards), and a propensity to ask their grad students to do far more manual labor than a computer scientist would like to see them do. I still remember talking to the poor grad student who studied clams. She was to be away for the summer and had hired an undergrad to help with feeding the clams; they all died, so this meant that she had another nine months of grad school ahead of her. To actually have physical specimens was a whole different world than what I was used to. You think: *"We can help you! You need a computer in that loop, but you also need to keep the human in that loop; you're not going to automate away the problem with curing cancer anytime soon."* It's an extremely complex multi-layer thing, a perfect design target for visualization.

The systems biology world is also interesting; I think of my collaborations with the Harvard-MIT Broad Institute and especially with Angela DePace [18], who worked on genomics and had fruit flies as her model organism. She had these rich and complex computational pipelines that we were trying to help support. It was an interesting example of a general principle, where any time you have a parameter, you should ask *"What if I picked a different number?"*; if you have something set to 0.7, how do you know it shouldn't be 0.8 or 0.6, how can you explore that? Complex computational pipelines are really rich territories for building visualization and visual analytics interventions that could help people understand which of their assumptions are correct, or what surprising things are there that they hadn't thought of.

I've also worked with microbiologists like Bob Hancock. Our Cerebral project [1] was with Agilent, another industrial collaboration. That one was interesting because it was funded by a company whose goal was generating open source software through the Cytoscape package. A lot of our software has been distributed as open source, and it's something I highly value.

Matthew: I'm glad that you pivoted the question to what was most interesting or revealing. The flip question is whether there is a research project or a paper that you've worked on that you feel is under-appreciated or under-circulated within the visualization community. Is there a paper that wasn't published at a place like IEEE VIS or EuroVis, but nevertheless is something that the visualization community should be more aware of?

Tamara: I have a couple of favorite underdog projects. One was published in the *Information Visualization* journal with Jessica Dawson as first author: *A Search-Set Model of Path Tracing in Graphs* [11]. It's a much longer paper than we usually write, which is part of why we decided that it had to go straight to a journal; there was no way we could chop it up into conference paper sized chunks that would make sense. It started with a qualitative investigation of what people actually do in when tracing paths through node-link graphs. We did an intensive analysis where we got people to move their stylus along a path to understand exactly what they're doing. We built a behavioral model based on insights from that. We then tried to instantiate the model, to do tests with multiple regressions to really understand not only if this significant, but to determine which of the factors are really doing the same thing versus which ones are doing different things. It is a hefty paper, I admit, and it hasn't gotten as much love as some of the other papers because it didn't have a conference presentation.

Another underdog was our *Dimensionality Reduction in the Wild* project [27]. It got published at the BELIV workshop [6], but it was a long path to get it out there. This was a qualitative field study with a few dozen high-dimensional data analysts. We were trying to understand how people use these sophisticated algorithmic tools and whether the assumptions at the beginning of all the algorithms papers, which assume a smooth and densely sampled manifold, and I still remember the first time I read that, I thought *"How could you possibly assume that? Why would that be true? Most of the* *time it's completely false. Does this algorithm even still work?"* We would ask: *"why are they using the Swiss Roll dataset as as their benchmark dataset? It's got a crazy structure that has nothing like what we typically see in visualization."* Ten years after I first started asking these questions, I felt like that we had finally shed light on what people were doing. We discovered that some people care about dimensions, some care about clusters, and some care about both. Untangling that was the key to understanding why people use certain benchmarks and language.

Matthew: I'm glad you mentioned that project, since dimensionality reduction factors into another domain that you've undertaken problem-driven research in, being journalism. Specifically, I'm referring to applying dimensionality reduction to large document collections, such as what a journalist might get from a freedom of information request or from a whistleblower leak [4].

In a 2008 panel talk [20], you highlighted *total political transparency* as a grand challenge for information visualization, and to this end, we've seen how journalists use data and visualization to hold political entities to account. You've engaged with the journalism community several times since then, with a keynote address at the 2016 Computation + Journalism Symposium [23] and a string of collaborations with journalists [4, 12, 13]. Obviously, the landscapes of visualization and journalism have changed a lot over the past decade. Reflecting on this now, how can these communities work together to address the challenge that you posed?

Tamara: For a number of years, I saw the visual analytics community really focus on intelligence analysis as a target domain – a lot of their funding came out of homeland security, at least in the United States. I wanted to make sure that as a community, we didn't just focus on intelligence analysis, but to make sure that we had transparency, openness, and sharing of information. I have thought of journalism as a really important force for good, and I think that continues to be true. Even though visual analytics has become much more integrated into the visualization community and is no longer as dependent on homeland security funding, it left a tendency in the field to favor intelligence analysts over journalists. I wanted to go in another direction.

We worked with hacker-journalist Jonathan Stray on the Overview project [4], which had to do with analyzing large text document collections. Then came TimeLineCurator with Johanna Fulda [12], in which we tried to speed up the process of creating timelines for news stories. Most recently, I worked with Steve Kasica and Charles Barrett on Table Scraps [13], which looked at data wrangling in a journalistic context.

Journalism is a really interesting domain because in some sense, it's so broad: *"What do you write stories about? ... Well, almost anything."* At the intersection of two broad topics (computer science and journalism), and especially the intersection of visualization and journalism, there's a lot of energy there. One the reasons that D3.js [2] got to be the force that it is today is because Mike Bostock developed it and then immediately went to go work for *The New York Times*, where he popularized doing really sophisticated things with visualization. At an early VIS conference, there was a capstone from someone who worked at *The New York Times* who said *"scatterplots are too complex for our audience."* Not even ten years later, we had *The NYT*'s Amanda Cox [8] giving us examples of incredibly sophisticated things that they were doing that went far beyond scatterplots. These days, data journalism is a force for educating people about what's possible, providing visual literacy as part of their narrative mission to inform people.

Matthew: Beyond the grand challenge of political transparency and how it manifests in data journalism, what are other grand challenges for visualization in the 2020s?

Tamara: It's a good question. I never have a quite as coherent of an answer as I wish I did. When people talk about grand challenges in many fields, they're talking about specific things: get a rocket to the moon or cure cancer. These are easy-to-articulate grand challenges. Visualization is a situation like statistics, where a lot of what you're doing is facilitating the goals of others, so it's harder to articulate grand challenges.

If a grand challenge is saying *"here's a city that I'm trying to get to,"* what visualization often has to offer are methods for building roads, and you could use these to build roads to a lot of different places. It doesn't just have to be that one city over there; if you learn how to build a road to that city, we're going to be able to build more and better roads to other cities as well. Using this extended metaphor, my goal is to make roads of *human-in-the-loop visual data analysis*. I want to make it easier and better to build these roads. It's hard for me to articulate that as a specific grand challenge. We want to continue with visualization being something that is powerful, flexible, and even more pervasive than it is now. It's not exactly a grand challenge, but it's what I want to spend the next couple of decades on.

Matthew: Your response brings to mind your methodological and theoretical work. You've been bringing lessons from human-centered design and human-computer interaction (HCI) into the world of visualization. I think especially of your nested model paper [21], your design study methodology paper [28], or our task typology paper [5]. What else can visualization learn from HCI and adjacent fields?

Tamara: There continues to be a lot that we can learn from HCI, but also we have things that we can bring to HCI; the flow is two ways. It's ironic because when I was getting my PhD, I didn't think of myself as an HCI person at all; I was in a graphics group and was focused on algorithmic work. I started to wonder: *"is this helping?"* I started getting into lab studies to answer the question of whether my proposed algorithms and techniques might actually be working. I still didn't know the answer, so I started getting into qualitative field work and more ethnographic methods. I still don't know the answer to the questions of *"what do they need?"* and *"am I building the right thing?"* For years, I've thought that maybe the answers are out there and if I just read the right HCI papers, I would find the answers. I found them to be surprisingly elusive, even as HCI papers remained useful. You have to take what's known in HCI papers and glean from them what's relevant to visualization. There's a lot in HCI that we can learn from, and there's a lot that doesn't directly transfer. For example, Computer-Supported Cooperative Work (CSCW) addresses the squishy problems of figuring out what people need and how to technologically intervene with better tools that help people get their work done, but CSCW doesn't typically deal with the volume of data that we have in visualization. We've also got a lot of these squishy humans. But we also face data science problems: a need for mathematical and algorithmic foundations to have all this to work with real data. It's that intersection that makes visualization really interesting.

Visualization is adjacent to so many fields. I've been doing it for (amazingly enough) a number of decades and there's still so much to read, so much to learn. A lot of my theoretical papers and my book [22] were all written to scratch my own itch, to figure something out for myself and write it down. Writing down helps one think and has this nice side benefit that everyone else can read them.

Matthew: I'm glad you mentioned your book; it's certainly one of the most notable achievements in visualization of the past decade. Reflecting now on the process of writing the book and the impact that it's had since its publication, how has the book changed the way you teach or conduct research?

Tamara: The book has changed how I teach. I now have this thing I can just ask people to read, or I can use as the basis of lectures. Before I had the book, I would assign (in retrospect) a somewhat harsh number of papers for the grad class to read every week. The point I wanted them to get out of a paper was not necessarily what you would get as the naïve observer. There's what I saw in the paper versus the words on the page, and they could often be quite disjoint. There were all these principles that I wanted to be able to articulate. I would say how a paper was an example of that principle, but it wouldn't actually state the principle clearly in the way that I wanted to. I'm now able to cover more territory more deeply, but with less reading time for the students, with a book that actually draws all of these ideas together in one spot.

What I didn't understand was how much the book would help my research. For one thing, I agonized about how to structure the book. What was really unsatisfying to me was that everyone's visualization class talked about a bunch of techniques, and it felt like a grab bag. At the end, they'd say *"We're gonna have a lecture on evaluation . . . you could do this, or you could do that."* I found it unsatisfying that techniques and evaluation were treated separately. I wanted to integrate them and I spent six months writing the nested model paper [21] as a way to outline the book, to figure out how to structure things. I needed to write this paper to start the book, and then our abstract task typology paper [5] was the other research that I needed to end the book. I had tried to write a paper about tasks back in 1999, when I was still in grad school, and after three months I thought *"I don't know enough to write this; I'm won't know enough to write this paper for a decade."* At the time, a decade sounded like a long time; but then I blinked and more than a decade had gone by. This time, you had jumped into the fray with a bunch of amazing ideas and enthusiasm. I didn't want to write the paper originally, or to write a book chapter about tasks. However, everyone who test drove the book remarked on there being a missing chapter on tasks. So we figured out how to think about tasks in an abstract way, which I still think is a really hard and interesting problem. Our paper became the basis for the chapter that I needed to finish the book. Even though the chapter appears early in the book, it was added late in the writing process.

Matthew: After publishing your own book, you've also become the editor of an entire book series on visualization, with each book synthesizing a different area of knowledge in visualization. Where else do you see a need for synthesis? In other words, what yet-to-be-written books about visualization do you want to read?

Tamara: There are lots of them, and luckily some of them are even getting worked on. I'm always on the lookout for more. People should talk to me if they have even the faintest glimmering of an idea of a book in their head. There is still very much a need for textbooks that cover a substantial fraction of the whole field, because they come from different points of view.

There's also a place for books that are more focused on particular topics. There's a desperate need for books where biology and genomics meet visualization. The problem is so broad: you take the cross product of biology and visualization, and you get a lot of things. It could be even more targeted than that, such as a book on omics visualization. There are many fields where we need more depth: domain-oriented visualization books for journalism, genomics, finance, business intelligence, or the digital humanities. All of these areas could benefit from more books.

I also think of visualization books about specific types of data. There was Richard Brath's recent book on text visualization [3], but there's still a lot more territory to be covered. There's machine learning, where if you think of the models themselves as a data type, you can try to visualize them, and that's a very rich and hot area that would be great for a book. It's an area where it's hard to keep up, because the work is moving so fast, but books give you an opportunity to synthesize.

Being a methods geek, there's of course a need for more books about methods, such as a book about how to do qualitative work that is specific to visualization, hitting it from different angles. Rather than take general methods from the social sciences or from HCI, this book would focus on how to do qualitative analysis or ethnography in a visualization context. For example, like guerilla usability, we need guerrilla ethnography, or what would actually work in a practical context if you don't have infinite time.

All of these books could make the world a better place. Hopefully people will write them.

Matthew: The need for more synthesis through books could also be a reflection of how the visualization community has become more professionalized over the past decade, such as with the emergence of visualization-specific roles within organizations. Has this shift affected your teaching or research collaborations?

Tamara: It hasn't yet changed with whom I collaborate, because I think a lot of the interesting stuff is being done by freelancers or small boutique visualization design firms that are not likely to be in a place for academic collaboration. But this shift informs the kinds of things that I teach. There is now so much more demand for visualization practitioners; it's why it makes sense for me to offer a class at the undergraduate level. In my *'next steps'* lecture, I love showcasing the *Data Sketches* project by Nadieh Bremer and Shirley Wu [7] or the *Dear Data* project by Giorgia Lupi and Stephanie Posavec [15] as examples of how you can do things as a practitioner that are utterly and completely out-of-the-box; it's not just a bunch of bar charts and pie charts.

There's three ways that I have become more aware of what's happening in the world of practitioners. One is the *Data Stories* [29] podcast from Moritz Stefaner and Enrico Bertini. Another is the *VisInPractice* event at the IEEE VIS conference. The third is datavis Twitter; there are some academics that are vocal, but even more practitioners, and that's how I hear about a lot of things. One of the active voices there is Alberto Cairo, who straddles practitioner and researcher roles as somebody who's in both journalism and visualization. I'm super delighted that he's joined me as a co-editor of the AK Peters Visualization series. We're hoping to get more practitioner-oriented books, and we have quite a few in the pipeline, including the *Data Sketches* book [7] that just came out.

Matthew: Both you and Alberto are also on the board of advisors for the Data Visualization Society (DVS) [10].

Tamara: Yeah, which is another incredible phenomenon that went from zero to sixty in ten seconds. Elijah Meeks, Amy Cesal, and Mollie Pettit had this idea of of getting people together that care about visualization. Then they blinked and ten thousand people were members. Now it's up to seventeen thousand; it's got energy, passion, and geographic breadth. A lot of people involved in the DVS are not necessarily early in their personal careers, but they might be early in terms of their engagement with visualization. It's acting as this really fantastic on-ramp from nothing to something, and from something to something deeper. The DVS helps people get to the next level, no matter where they're at. I'm very grateful; when I started in visualization, it was a much more academic world, a niche. Watching this community blossom has been pretty amazing.

Matthew: Looking now to the future, where do you see your research going in the next few years?

Tamara: For one, I'm interested in human-in-the-loop curation and visual analysis. It's a bit of a cliché to say that I'm interested in the intersection of machine learning and visualization, because everyone and their pet hamster is interested in that; I am too. I see two obvious directions. The first is *visualization for machine learning*, such as visualizing what's going on within a model training epoch. Another direction, which is even more interesting and in some ways more difficult, is *machine learning for visualization*, or how do you truly synthesize automation and human-in-the-loop control for some third task: not to do machine learning itself, but to use it to do something else.

I also continue to be interested in biology, genomics, medicine, and health. Over the past few years, I worked with Anamaria Crisan, who's now at Tableau, and she has deep knowledge in the domain of genomic epidemiology. I learned a lot from her about public health, a topic that has taken the world by storm as a result of the pandemic. It's incredibly difficult to actually use visualization effectively in a health care context; it's so heavily regulated. You have so much to do to get anything adopted in a clinical or public health setting. There's a lot that could be done, but also a lot of pitfalls and quagmires to be careful of. The data is extremely sensitive. Moreover, there's methodological challenges in adapting human-centered methods to a public health context; Anamaria was first author on a paper [9] about this that affected how I think. Going forward, I continue to be interested in the prospect of visualization design studies in this domain.

Matthew: You've mentioned several of your students and post-docs throughout this conversation. Alumni from your group have gone on to successful careers in academia and industrial research labs, or they're out in the world building products. What's your secret in terms of mentorship?

Tamara: I don't know if it's a secret, exactly. When I started this professor job, I thought that I would have a policy of how I interacted with students. That idea quickly went out the window because everyone's different. There's no one-size-fits-all policy. The conversations I have with trainees are always different. The conversations I have with a new Master's student who's learning the basics (such as how to do a literature search) is totally different than conversations I have with a post-doc who has questions about how one gets a faculty job at a desirable place. Every student I work with takes the same amount of time, but the content of the conversations varies dramatically. Maybe they're coming in with ten years experience in medicine, or maybe they're coming in with a fine arts degree, or a background in astronomy or cognitive science. Maybe they are incredibly organized and have mastered time management, or maybe that's something they need to learn. Maybe they're an amazing hacker, or maybe they really know HCI research methods. They all come from a different place.

My job is to help trainees figure out what they want to be and make sure that they're getting prepared for that during their time with me, so they get catapulted on the trajectory of their dreams. That's both the amazing and terrible thing about being a professor: nobody stays. You know from day one that it will be temporary. You will be a part of their lives for a while, and then they will go on to something else. There's a cub scouts motto about camp sites, which is *leave it better than you found it*; I feel like that's your job as a professor with trainees, to help them become what they want to be during their time with you.

Matthew: Along any trainee's trajectory, there's bound to some failure and rejection, especially early on in a research career. Do you have any advice for junior researchers about this?

Tamara: I have pragmatic advice. In cases where reviews are negative, or even if your work gets accepted but reviewers have misunderstood something, you should read the reviews, swear (a lot), don't write any email, don't do anything, sleep on it a few days, and come back to it when you're calmer. The way I frame it is not *"That evil reviewer they didn't take my paper!"* or *"what a dumb reviewer; they didn't get my point."* Instead, think about how you can write the paper in a way that nobody could miss the point. If a reviewer could miss the point, then maybe a lot of readers could too. It's an opportunity to do a better job of explaining what is cool about your project. Obviously, rejection hurts. It still hurts now, and I've certainly seen a lot of it, but it is also something you can really learn from. Many first-time authors might not know how to deal with a major revision or a rebuttal. They might not understand just how much work goes into writing responses to reviewers' concerns, because first-time authors never see these letters; they're used to reading papers. Hopefully, they've read a lot of papers, but they haven't read a lot of responses to reviewers. How to write and structure those is important to learn early in your career. It's certainly something I give a lot of guidance on, because I have written so many of them.

The other slightly roundabout answer to your question is to try to structure your work to make it less likely that it is rejected in the first place. This means no last-minute rushing a meet a deadline. In my group, we have a somewhat unusual way of doing things. Three months before a paper deadline, the project lead prepares a *pre-paper talk*, a slide deck that becomes the outline of the paper. We imagine as if the paper has gotten accepted and it was being presenting at the conference. This is rapid prototyping for papers through slide decks instead of through prose. It's faster, and we iterate on that slide deck a lot before it gets presented to the whole research group. The group takes out their machetes, slashes it to bits, and argues about it. Someone might say *"Slide #54 should have been slide #3 . . . I didn't understand the main point of the paper until you were almost done!"* or *"You left out this completely crucial idea!"* There's often some shared unspoken assumptions among the co-authors that the presenter forgets to say explicitly. This helps us workshop the paper's message at a slide-by-slide and section-by-section level. It helps us to identify the research contributions, which can evolve dramatically from what you thought they would be when the project started.

Early career researchers may think that you start a research project knowing what you're doing. In reality, you often don't start a project knowing what you're doing; you can end up doing something completely different, and you have to retroactively figure out what you did, to frame it in a way that would make anyone else care. We spend a lot of time on that, on writing as discovery. Between this and the pre-paper talks, we end up with for papers that are, on average, stronger than a typical paper. We put a lot of time and effort into trying to cut down the chances that they'll get rejected. Having said that, alas, not all of our papers get accepted; we are not immune from paper rejections by any means whatsoever. However, we have a process where if it's clear that the paper really isn't ready, we slip it to the next deadline rather than undertake a death march to get it submitted. I just don't think that's a good use of time, so we typically try to avoid that.

Matthew: Earlier you mentioned meeting trainees where they're at, acknowledging their varied backgrounds and interests. Regarding your own background and interests, beyond those directly related to your research, what have you found to be helpful?

Tamara: I have a lot of interests that don't directly or obviously affect my work but make me happy as a human being. It's important to have some balance. I find it a bit hypocritical to use the term *work-life balance*, because I think mine is a bit abominable and I could do much better at it. But you should aspire to it and do a reasonable job of it. It gives you energy and joy in the world. I do think you should spend time on things that don't directly give you a set of skills that you'll apply to work immediately; you should just do things because you enjoy them. I read an incredible amount of novels because it makes me happy, not because it helps my work, but just because that's how I get through the world. I read a lot of science fiction and fantasy, but I think the extent to which this reading informs my work is frankly, pretty minimal. But if I didn't do it, I don't think I could do my work.

Matthew: So you don't encounter ideas for visualization projects inspired by the fiction you read?

Tamara: Less than you might think. While I was writing my thesis between 1995 and 2000, I thought about having a little appendix on visualization in science fiction, critiquing strong or not-so-strong examples. By 2000, I decided not to take that on. It would have taken too long and there was too much out there. It's difficult to engage piece-by-piece with individual science fiction novels

Science fiction intersects with computer science in some really weird ways. There are certain parts of computer science in which science fiction can be wildly optimistic. In other areas, it's wildly pessimistic. Writers often posit way beyond what we have now, such as artificial general intelligence. With respect to visualization, many writers seem to be five years behind what's current now, even though their work is set 300 years from now, and you think *"My God, I hope HCI has gotten further than that by the 24th century!"* A lot of science fiction tends to lag in terms of HCI considerations. In many science fiction novels, authors don't describe the visualization in sufficient detail for you to tell whether it's a good or bad idea. There's very few science fiction authors who actually know enough computer science to have interesting things to say; exceptions that come to mind are Neal Stephenson, Charles Stross, and Cory Doctorow, but none of these focus much on visualization. Although I have to admit that Stephenson's *Snow Crash* [30] inspired a whole generation of people to implement a virtual reality modeling language based on it. Of course, this proved tricky: the speed of light is an issue when you're dealing with VR environments that are supposed to be global around the world. It was a novel, it's not a specification. Some people got a bit confused about that.

Matthew: You settled in Vancouver nearly two decades ago, and you've become a part of the Canadian visualization community. Are there particular challenges that are unique to Canada that this community can address?

Tamara: I'm not sure if I can identify challenges that are uniquely Canadian. However, I have seen a real growth and vibrancy of the Canadian visualization ecosystem. Western Canada in particular is just a hotbed of cool stuff. It's exciting that Sheilagh Carpendale jumped ship to Simon Fraser University, where she joins Lyn Bartram and John Dill; we also have Charles Perin at the University Victoria and Wesley Willett at Calgary.

Vancouver is also part of the larger Cascadia Corridor, where there's so much going on. Down in Seattle there's Tableau Research, Jeff Heer's group at the University of Washington, and Microsoft Research, as well as smaller operations like Danyel Fisher at Honeycomb. There's also newer people coming in with energy and enthusiasm, which is really fun and exciting.

Matthew: Before we wrap up, do you have any acknowledgments, thanks, or shout-outs that you'd like to extend to anyone who has helped you throughout your career?

Tamara: There's so many people to think. At the Geometry Center, I got amazing mentorship from Charlie Gunn and Stuart Levy; I also benefited a lot from conversations with Mark Phillips. Another strong figure for mentoring me at this time was Delle Maxwell; although she didn't formally work at the Center, she was the art director on all of our videos. I learned a huge amount from her. Words can't completely express what I owe to Pat Hanrahan for his mentorship.

There's also a cavalcade of people who mentored me who have gradually shifted from mentors into friends and colleagues. I think of Maureen Stone who's now at Tableau Research, but I originally started interacting with her when she was hanging out at Stanford. I also think of Betsy Zeller, I first met at SGI.

I don't know if you would call them all mentors, but everyone that I've ever worked with has taught me something, whether they be co-authors or my students. Every mentorship is a two-way street, where both people get something out of it. I have learned an immense amount from all of the students I've worked with over the years, and I'm incredibly grateful to them. I feel like I'm surrounded by amazing mentors.

Matthew: Thank you so much for speaking with me today. Congratulations again on this award and I look forward to your talk.

Tamara: Thank you so much. It's been fabulous to speak with you.

REFERENCES

[1] A. Barsky, J. L. Gardy, R. E. Hancock, and T. Munzner. Cerebral: a cytoscape plugin for layout of and interaction with biological networks using subcellular localization annotation. *Bioinformatics*, 23(8), 2007. http://www.cs.ubc.ca/labs/imager/tr/2007/barskya_cerebral_appnote/.

[2] M. Bostock, V. Ogievetsky, and J. Heer. D^3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of InfoVis)*, 17(12), 2011. https://d3js.org/.

[3] R. Brath. *Visualizing with Text*. CRC Press, 2020. https://www.routledge.com/Visualizing-with-Text/Brath/p/book/9780367259266.

[4] M. Brehmer, S. Ingram, J. Stray, and T. Munzner. Overview: The design, adoption, and analysis of a visual document mining tool for investigative journalists. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of InfoVis)*, 20(12), 2014. http://www.cs.ubc.ca/labs/imager/tr/2014/Overview/.

[5] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of InfoVis)*, 19(12), 2013. http://www.cs.ubc.ca/labs/imager/tr/2013/MultiLevelTaskTypology/.

[6] M. Brehmer, M. Sedlmair, S. Ingram, and T. Munzner. Visualizing dimensionally-reduced data: Interviews with analysts and a characterization of task sequences. In *Proceedings of the Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization (BELIV)*, 2014. http://www.cs.ubc.ca/labs/imager/tr/2014/DRVisTasks/.

[7] N. Bremer and S. Wu. *Data Sketches*. CRC Press, 2021. https://www.datasketch.es/.

[8] A. Cox. How editing and design changes news graphics, 2011. IEEE VisWeek capstone address: https://doi.ieeecomputersociety.org/10.1109/TVCG.2011.256.

[9] A. Crisan, J. L. Gardy, and T. Munzner. On regulatory and organizational constraints in visualization design and evaluation. In *Proceedings of the Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization (BELIV)*, 2016. http://www.cs.ubc.ca/labs/imager/tr/2016/regulatory/.

[10] The Data Visualization Society, 2021. https://datavisualizationsociety.com.

[11] J. Q. Dawson, T. Munzner, and J. McGrenere. A search-set model of path tracing in graphs. *Information Visualization*, 14(4), 2015. http://www.cs.ubc.ca/labs/imager/tr/2014/SearchSet.

[12] J. Fulda, M. Brehmer, and T. Munzner. TimeLineCurator: Interactive authoring of visual timelines from unstructured text. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of VAST)*, 22(1), 2016. http://www.cs.ubc.ca/group/infovis/software/TimeLineCurator/.

[13] S. Kasica, C. Berret, and T. Munzner. Table scraps: An actionable framework for multi-table data wrangling from an artifact study of computational journalism. *IEEE Transactions on Visualization*

and Computer Graphics (Proceedings of InfoVis), 27(2), 2021. http://www.cs.ubc.ca/group/infovis/pubs/2020/table-scraps.

[14] S. Levy, D. Maxwell, and T. Munzner. *Outside In*, 1994. https://youtu.be/4cCumubgDiA.

[15] G. Lupi and S. Posavec. *Dear Data*. Princeton Architectural Press, 2016. http://www.dear-data.com/.

[16] D. Maxwell and T. Munzner. *The Shape of Space*, 1995. https://youtu.be/-gLNlC_hQ3M.

[17] P. McLachlan, T. Munzner, E. Koutsofios, and S. North. LiveRAC: interactive visual exploration of system management time-series data. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, 2008. http://www.cs.ubc.ca/labs/imager/tr/2008/liverac/.

[18] M. Meyer, T. Munzner, A. DePace, and H. Pfister. MulteeSum: a tool for comparative spatial and temporal gene expression data. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of InfoVis)*, 16(6), 2010. http://www.multeesum.org/.

[19] T. Munzner. H3: Laying out large directed graphs in 3d hyperbolic space. In *Proceedings of the IEEE Visualization Conference, Information Visualization Symposium and Parallel Rendering Symposium (VIZ)*, 1997. http://graphics.stanford.edu/papers/h3cga.

[20] T. Munzner. Outward and inward grand challenges, 2008. IEEE VisWeek Panel: Grand Challenges for Information Visualization; https://cs.ubc.ca/~tmm/talks/vis08/vis08.pdf.

[21] T. Munzner. A Nested Model for Visualization Design and Validation. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of InfoVis)*, 15(6), 2009. http://www.cs.ubc.ca/labs/imager/tr/2009/NestedModel/.

[22] T. Munzner. *Visualization Analysis and Design*. CRC Press, 2014. http://www.cs.ubc.ca/~tmm/vadbook/.

[23] T. Munzner. Visualization and journalism: Four vignettes, 2016. Computation + Journalism Symposium keynote: https://youtu.be/UKihJzHqipc.

[24] T. Munzner, F. Guimbretiere, and G. Robertson. Constellation: A visualization tool for linguistic queries from MindNet. In *Proceedings IEEE Symposium on Information Visualization (InfoVis)*, 1999. http://graphics.stanford.edu/papers/const.

[25] M. Oppermann, R. Kincaid, and T. Munzner. VizCommender: Computing text-based similarity in visualization repositories for content-based recommendations. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of InfoVis)*, 2020. http://www.cs.ubc.ca/group/infovis/pubs/2020/vizcommender.

[26] M. Phillips, S. Levy, and T. Munzner. Geomview: An interactive geometry viewer. *Notices of the American Mathematical Society*, 40(8), 1993.

[27] M. Sedlmair, M. Brehmer, S. Ingram, and T. Munzner. Dimensionality reduction in the wild: Gaps and guidance. Technical report, University of British Columbia Department of Computer Science, 2012. http://www.cs.ubc.ca/cgi-bin/tr/2012/TR-2012-03.

[28] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of InfoVis)*, 18(12), 2012. http://www.cs.ubc.ca/labs/imager/tr/2012/dsm/.

[29] M. Stefaner and E. Bertini. Data Stories, 2012. Podcast. https://datastori.es/.

[30] N. Stephenson. *Snow Crash*. Bantam, 1992.