

A conversation with CHCCS 2019 Achievement Award recipient Dr. Carl Gutwin

Dr. Carl Gutwin
The Interaction Lab, University of Saskatchewan

ABSTRACT

A 2018 CHCCS Achievement Award from the Canadian Human-Computer Communications Society is presented to Dr. Carl Gutwin for his many contributions to the field of human-computer interaction (HCI) research. We especially note his impactful work on groupware such as inter-personal awareness and the mechanics of collaboration. His work spans the breadth of HCI from very technical systems papers, to inventing new interaction techniques, to exploring the application of social theory to technology design and to his extensive care and rigor in evaluative methods. 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 are the edited responses from Carl Gutwin to questions posed by Sheelagh Carpendale that took place during March 2019.

THE INTERVIEW

CHCCS: Hello Carl. Congratulations! I am delighted that you are receiving this CHCCS achievement award!

Thank you very much! It's a real honour to be among the luminaries who have won this award in the past.

CHCCS: I believe you have been at University of Saskatchewan 20 years, now? Looking back, would you say that a lot of things have changed, at your university and in general with Canadian research in human computer interaction?

Wow, 20 years. It's interesting, this idea of looking back, because a lot of the time it feels like I've only just started, and so how could anything have changed? I can't believe that 20 years has gone by – a statement that will seem completely weird to new professors, and laughably familiar to others, I'm sure.

But, if we're looking back (and I had to consult Wikipedia to find out what was happening in the world in 1998 when I started), here are some somewhat-random observations. In the first part of 1998 I was doing a postdoc with Ian Witten, my grand-supervisor, at Waikato University in New Zealand, learning about text processing and information theory and digital libraries. That was a great experience – always good to do something completely different than what you did in your PhD – and led to a paper on keyphrase extraction that has turned out to be very popular. I also remember when I heard that Google was going public, and I thought “Who on earth would buy shares in a company that gives search away for free?”. So of course I didn't buy any shares. Which is why I'm in academia and not business.

In the past 20 years, user interfaces (on the desktop, anyway) have changed very little – just a bit less ugly compared to Windows98. One of the big HCI revolutions, of course, was the iPod touch and the iPhone – when was that, 2007? I remember that when the iPhone came out, thinking “but we've had handheld PDAs for years, and Windows CE devices even have CDPD

modems! Who would want one of these?”. So I didn’t buy any shares in Apple, either. Somehow I missed the importance of doing the design right, and having a Svengali to do your product announcements.

In my PhD I studied shared-workspace groupware, and now, 20 years later, I’m studying...shared workspace groupware. But on the web! So it’s totally different. The CSCW community has expanded a lot from what we were working on in the 1990s, doing Facebook studies, Twitter studies, Twitch studies...but I still like shared workspace groupware. This is the good thing about just about any research area – there’s a lot of work to keep you interested over the long term. And of course I’ve done a lot of different things as well along the way.

CHCCS: Do you have suggestions about strategic areas for new HCI faculty?

Good question – there are lots of interesting new areas! But let me talk for a moment about taking some of the ideas that already exist and applying them to the contexts that we’re looking at today. I don’t mean for this to sound like a “we did everything in the 1980s” kind of answer; the important thing is that some of the techniques that we already know about become completely different when the datasets are as big as they are now. An example – I think that there’s a lot to be done in what I’m starting to call “the human factors of big data”. Ideas that people were working on in the 80’s and 90’s, like one of my favourites – edit wear and read wear – have become even more applicable now that users are trying to deal with datasets of millions of items rather than tens or hundreds of items. Navigation, revisitation, awareness, all of those things were already interesting problems back when I learned about them in the early 90’s, and although not much research is being done on these problems anymore, at least in terms of seeing the ideas in production systems, the scale of the problem has exploded. The wheat genome, something I’ve been using for visualization, has 300,000 genes, and 17 billion base pairs of DNA. Try navigating in that dataset without any revisitation support!

The flip side of this question is that any researcher will always miss many of the strategic areas that will be important for HCI in the next twenty years – you already know how good I am at guessing which technology companies will be successful! I’m pretty sure that VR is going to be a big deal, but when I put on a VR headset I get sick in about 30 seconds, so some potentially-interesting topics are off the table for me.

CHCCS: On the topic of previous work, what would you consider to be your most significant contribution over these years, or perhaps your favorite result?

First, I think we shouldn’t worry too much about whether our work will be influential. Often in HCI we’re doing things that make life a bit better for users, and that move the field a little bit – but that’s just science, of course, and I’m fine with that. I think we can be happy to be part of that larger enterprise even if each individual result isn’t revolutionary. But of course this can still be a question about what I personally am proud of, and I guess there’s a couple of things. First, the awareness work has been influential, and I’m still very fond of those papers, although perhaps because I remember all the late nights in Calgary working on those systems and experiments and papers. Second, the spatial-memory work that I did (and am still doing) with Andy Cockburn has been extremely interesting, and I think it’s an important idea that I’d like to see made more use of in real-world interfaces. And third, what I really enjoy doing is building application-level networking techniques for real-time distributed groupware, which is definitely a niche area and not one that a lot of people tend to read, so not very influential overall (if any students are reading this whose eyes just lit up, please contact me). Last, I’m also proud of trying to do a good job of writing all of my papers – I love getting an Introduction to the point where it’s clear and interesting and tells a good story. I hope that students who read my papers are inspired to become better writers.

CHCCS: How do you choose research topics and directions? Do you think that is a role for chance in research?

The most important thing for me over the last twenty years has been working with great collaborators – that is the best way to find new topics, learn new things, and have fun while you’re doing it. I have been lucky to have collaborators who are fantastic researchers and wonderful people – Saul Greenberg, Ian Witten, Regan Mandryk, Nick Graham, Pourang Irani, you too Sheelagh, and my main collaborator of the last 20 years, Andy Cockburn. Work with great people and listen to what they want to do (and say “yes – let’s do that!”). Students also bring you new questions that you otherwise never would have looked at. I remember that Miguel Nacenta came to me and said that he wanted to change the cursor according to the user’s perspective. This was a great idea – but of course, the first time he said it I had no idea what he was talking about, and he had to keep coming back to me about ten times before I understood the implication of the idea. And lately, I’ve been working with genomic scientists and plant breeders, and they definitely have to tell me multiple times what they’re talking about (synteny over phylogeny – the next big thing!). So let me amend that advice – pick great collaborators who are nice people and who are really patient at explaining their great ideas to you!

In addition, I definitely end up with some research topics just by being easily distracted. A more complimentary way of saying that is that I'm interested in lots of things, and I can see HCI problems in just about anything. So, part of keeping things interesting for me is being able to work on whatever puzzle presents itself that month. Here's an example: I was sitting in Sylvain Malacria's CHI talk that was about encouraging the use of hotkeys, and Bill Buxton got up and said "but everyone uses touch devices now – and there aren't any hotkeys on a touch device." And I thought "wow, that's right – we should invent shortcuts for experts that work on touch devices." So we built FastTap, which was a fun project where we use multitouch chords as menu shortcuts, and it works really well, and it just grew out of that realization sitting in the audience at CHI. So chance plays a huge role if you're happy to be a bit of a dilettante.

With that said, if you keep doing research for long enough, you tend to go back to some of the same things, which is what I've done with topics like awareness and read wear. I kind of have an awareness filter and a read wear filter on my perception of the world, so I tend to see awareness problems and interaction-history problems wherever I look, and that means that over the years I have actually done some things that tie together across a longer term. Which is interesting to look back on – good thing we have Discovery grant proposals that force us to rationalize our last six years of research and show how it all fits together!

CHCCS: Are there strategies to problem solving or problem selection that you would pass on to others?

Empiricism is really powerful. Perhaps we know that already after two thousand years of science, but empiricism has had some bad press lately because of things like the replication crisis and p-hacking and so forth. The scientific method is still the most powerful way we have of discovering new things about the world. The flip side is that being honest in your empiricism is just as important as the method itself: science is a social enterprise, and lots of different aspects of science are based on the community trusting the researcher to do things correctly and honestly. This is a hard thing to do for junior faculty who have enormous pressure to get things published, but it turns out that being honest with yourself – for example, recognizing that an idea you've had is a dumb one – is also a valuable skill.

CHCCS: Do you think you have a most under-appreciated paper? Is there a hidden gem out there that got published at a less visible venue, but is something that people should be aware of?

Application-level networking for real-time distributed groupware, obviously :)

CHCCS: What parts of your previous work get the most use in industry? What do you think about tech transfer and commercialization?

Even though I've been producing stuff for 20 years, I think that it's still too early to answer that – it's hard to tell where things come from in software and devices and maybe things still haven't progressed to the point where my results can be practically applied. I hope that Google Docs knows about group awareness, whether through my papers or someone else's, and I hope that they're planning to add telepointers to their tools any day now. But I don't talk to those groups very often, so I'm not sure.

In terms of tech transfer, I'm still very much in favour of being able to carry out speculative and curiosity-driven research. Lots of the current funding opportunities are tied to industrial relevance, which is fine, but I'm glad that some of the programs still make it possible to do things that are relevant in other ways. Dr. P (Przemyslaw Prusinkiewicz) gave a talk at GI a few years back about how he discovered something completely novel about how plants form buds, just by being interested in trying to build a model of growth for one of his digital plant systems – so an important biological result came from just trying to figure out how it worked. That might not have happened if all of the research programs were tightly connected to industrial relevance.

CHCCS: IS there a relationship between quality of research and difficulty?

I don't think so – sometimes great results come easy, and minor results are really difficult to get. It's too difficult to define research quality anyway – look at the disparity between what papers win awards at CHI and what papers actually get cited over ten years. Having an insight about a problem might be super easy – right place, right time – and still be super important. The person who invented the sewing machine supposedly had the key idea in a dream, which doesn't sound very hard – so maybe we should all just sleep more so that the good ideas can just come to us.

CHCCS: Do you have any "behind the research" stories about results that were surprisingly more difficult than what you would have otherwise expected?

Everything is surprisingly more difficult than what I would have expected! But one place where this is always true is in trying to study groups doing work in groupware – like trying to quantify the value of an interface that provides awareness information. People are so unbelievably resourceful, that if you give them two interfaces, and one has great awareness displays that should support them in doing the task, and the other has none of that support, I am always astounded at how people can be fast and efficient and error free in the interface that is supposed to be terrible – and that everyone looking at the two interfaces would agree should be terrible! The almost-infinite reservoir of effort that people can apply to a task has made my empirical life much more difficult over the years.

CHCCS: Do you have comments/advice about work life balance?

Only that for academics, we have nobody but ourselves to blame for taking on too much. So, my advice is to learn to be better than I am at saying no (it helps to have a spouse who can sometimes say no for you). Also, for junior faculty, sorry, there is no such thing as work-life balance (but it gets better after a few years, so keep swimming).

CHCCS: Do you have comments about ‘failure’ and ‘rejection’ in research?

First, get used to rejection – it’s going to happen a lot, so try to figure out how to not take it too personally (which is hard to do, but gets easier after lots and lots of rejections!). Second, learn from each rejection. It isn’t always the case that reviewer two was an idiot. If we don’t make the brilliance of our work clear to the reviewers, how can we expect them to give us great reviews? This comes back to being a good writer – we can always improve the presentation and the way we tell the story. Almost always when I look at a rejected paper after six months so that I can revise it for the next deadline, I read it and think “What were we thinking? It’s no wonder the reviewers rejected this”.

Another story: I got one of my first independent papers rejected from CHI, and I realized that maybe I had no idea how to write a CHI paper. So I got one of Scott MacKenzie’s papers from the previous year’s conference, and I said, OK, what does he do in each section, and what does he do in each paragraph? So I rewrote my CHI reject using Scott’s paper as a template, and then it got accepted the next year (thanks, Scott!). So, use each rejection as an opportunity to figure out why you got rejected and what you need to do to improve for next time.

CHCCS: Are there are elements in your background that you found helped you in your career? Anything unexpected?

I’m glad that I got an English undergrad degree alongside my Computer Science degree – it forced me to become a good writer, and it made me read a lot of good writing. I actually applied to do an M.A. in English and got accepted at a couple of places, but chose to do Computer Science instead. So, the world missed out on having another starving English major – but I would have made a terrible English prof, because I’m a complete literalist and there’s not much room for empiricism in English Lit.

CHCCS: What might students like to know in preparing for successful careers in human computer interaction? What words of wisdom or advice would you give to new graduate students?

There are a million things! But if I have to pick one, let me go back to the issue of writing again, and encourage students to try and become better writers (which is really a life’s worth of work). Academic success is still based on written communication, and if you can write well, you have more chance of being successful. I tell my students that our goal is to write up the work well enough that if it gets rejected, it’s never because they didn’t understand it – it’s OK if they disagree with the value of the research, but it’s never OK that they reject something just because we didn’t clearly say what we were doing or why we thought it was important. Write well enough so that people can judge your work fairly – then all you have to do is pick good questions, and if you collaborate with great people, that’s easy!

CHCCS: Would you say there are interesting emerging directions that might influence future human computer interaction research? Where do you see the changing trends or exciting problems?

Prognostication, at least on the record, is kind of a fool’s game, but here are some guesses. The HCI of big data, of deep learning, of AI – whatever that means, it’s going to be important in the next few years. One aspect of this that’s going to be critical for HCI is interaction with automated and semi-autonomous systems – there are going to be self-driving cars all over the place in a few years, and at critical points (when they’re sliding off a Saskatchewan highway in January) they’re going to have to interact with the human. How should that work? I have no idea but we’d better start looking at it soon. Another great area is going to be interacting with environments where all of the flat surfaces are displays – this is based on Bill Buxton’s prediction

that displays are going to be sold like wallpaper (“a dollar a square foot”), but I think we have no idea at all what it will be like to interact with these kinds of environments. (Miguel Nacenta and his students are doing some great work in this area, by the way).

Last, I should highlight a recent trend that I think is important and is going to grow in importance over the next few years – the idea of thinking about what the possible negative implications are for the technologies that we invent. We’re great at thinking about the positive potential, but as we’ve all seen over the past few years, important rights like personal privacy are easily compromised by technology that was designed to do cool new things. Everyone focused on the positive, and we forgot about the negative. This goes back to being honest in the science that we do.

CHCCS: Do you think about looking ahead 50 years ahead?

I can’t wait! Every year, new unbelievably cool stuff keeps coming out – and building things with new technology is one of the things I like doing most. I will undoubtedly miss out on some of the latest and greatest – I used to believe that because I grew up with computers (unlike the generation before me) I would never be left behind by technology – but then Twitter happened, and I had absolutely no idea why anyone would want to do that. But the good thing is that there will be a lot of new territory to explore, and I can be an expert in whatever areas interest me, even if I don’t know what the latest thing is that the kids are using. Plus, in 50 years I bet we’ll have excellent application-level networking for real-time groupware systems!

CHCCS: Thank you very much for this conversation. You have given me lots of things to think about. Congratulations again on your CHCCS achievement award. I will look forward to seeing you at Graphics Interface!

Thanks very much for this talk, and for the honour of this award. See you at the conference!