Transcript: Conversation Design and Conversational AI Meets Conversation Analysis

[0:00:00]

Elizabeth Stokoe: Welcome to this National Centre for Research Methods In Conversation session entitled Conversation Design and Conversational AI Meets Conversation Analysis.

I'm Elizabeth Stokoe, Professor in the Department of Psychological and Behavioural Sciences at the London School of Economics and Political Science in the UK where I'm also Academic Director of Impact.

Also with me is Dr Saul Albert, Senior Lecturer in Social Sciences at Loughborough University UK, and Cathy Pearl, User Experience Lead on Gemini at Google and the author of the O'Reilly book, Designing Voice User Interfaces.

Saul and I have been research collaborators for several years now through our work in conversation analysis. For those listening who are unfamiliar, conversation analysis or CA is a field comprising six decades of cumulative science of social interaction. CA's origins are in sociology but it's spread its foundations across linguistics, psychology and communication, and its applied research has travelled across most academic disciplines from medicine to computer science.

CA is a method and tool for systematically analysing talk and embodied conduct across naturally occurring actual in the wild settings, but it's also a theory of interaction for understanding human sociality. CA doesn't study just talk, but all of the spoken and written embodied resources that people use to

interact, on the phone, in digital applications, in person, remotely and in many different languages.

In terms of the way we typically categorise social science research methods, CA isn't neatly qualitative or quantitative. It's perhaps both with the scope of its research reaching from studies that take apart single interactions to those that work with very large and often multiline data sets to investigate the universal nature of the machinery that drive social life.

One of the things I've been most interested in is the gap between what people say about how conversation works and how it is actually demonstrably and empirically organised. Despite myriad communication myths and stereotypes, conversation is actually leveraged for products of all kinds, from communication guidance, training and assessment tools to chatbots, voice assistants and the other conversational technologies that are today's focus. And it's this particular topic of conversational AI and technology that brings us all together.

Saul and I are going to talk through our collaboration by asking Cathy lots of questions, but first I'll invite them both to introduce themselves, first Saul and then Cathy, and then I'll kick off with the first of our questions to Cathy.

Saul Albert: Okay. So, I'll briefly introduce myself. My name is Saul Albert. As Liz described, I'm a senior lecturer in Communication and Media at Loughborough University. I am working on conversation analysis and conversational AI in that I'm looking mostly at how people with disabilities work in so-called smart home care settings. So, that's been the bulk of my empirical work, but I'm also interested in how do you use the methodologies and the theories of conversation analysis to understand this new domain of interaction? So, I'll hand over to Cathy at this point.

Cathy Pearl: Hi everyone. I'm Cathy Pearl. I've been working in the field of speech recognition and natural language for a couple of decades. Right now, as Liz

mentioned, I'm working at Google and I work on Gemini, formerly known as Bard, our large language model. Back to Liz.

Elizabeth Stokoe: Great. Okay, so, here's the first of our questions to you, Cathy, about this collaboration that the three of us have been working towards for the past few years. So, start off by telling us a little bit more about your role at the moment and maybe how you came to be working in this field in the first place.

[0:03:56]

Cathy Pearl: Yeah. So, it's been an interesting evolution. Right now, I'm working as a conversation designer, UX lead, and basically what that means is we think a lot about how to train the large language models, how to have a good conversation. I think a lot of people think, "Oh, large language models, they work out-of-the-box. You don't have to do any training or any thinking about it". But it turns out they also need training and work to really make them work properly and be able to converse with them in a way that is natural and makes sense. And how I came to this, I have been interested in talking to computers since I was a kid and had my own computer. I always wanted to have that experience and didn't really think it was possible when I was a kid watching movies like War Games and shows like Knight Rider. But then in graduate school I started taking some courses in human computer interaction, started to learn more about that, and then eventually I managed to get a job at a company called Nuance Communications, which is one of the first companies to use speech recognition technology for the public through automated phone systems. So, it's kind of how I got my start.

Saul Albert: Great. So, Cathy, maybe you could say a few words about how you first met Liz.

Cathy Pearl: Yes. So, this was back in May 2019 and I was giving a talk at the Google IO conference which is our annual developers conference. I was talking about conversation design and I had read Liz's book, I didn't know her, but I'd read her book, Talk: The Science of Conversation, and I had included something

I'd learned in the book about how questions can end either up or down. So, you might say, you know, “How are you?” Then you might say, “Do you want fries with that?” And that was something I was mentioning in the talk and little did I know that Liz was watching the talk and she tweeted about it. She didn't tweet about that, but she tweeted about the concept of repair. And I was so excited, like, “Dr Elizabeth Stokoe is tweeting about me. Oh, my gosh”.

So, I think after that, I reached out and thanked her. And then she suggested we meet up sometime. So, we did a video call because I'm in California and she's in England. And then next thing you know, Liz was convincing me to apply for a grant to work with her in Loughborough and the rest is history.

Elizabeth Stokoe: So, from your point of view, Cathy, what would you say is the goal of our collaboration? Because for those of you listening, the pandemic got a little bit in the way of our ability to meet in person, but we did eventually meet in person and the three of us spent a good week or more together a couple of summers ago. But yeah, from your point of view, what would you say it is that we are trying to do together?

Cathy Pearl: You know, when I read your book and I thought there's so much here that actually applies to my job day-to-day. And I think one of the things that's difficult as practitioners, we don't necessarily have a lot of time to read papers and really understand a field outside of our own domain because we're just like so, so, busy every day. But I thought, you know, what a shame and I tried to share some of the knowledge that I found from your book with my coworkers and I realised there's this wealth of information in conversation analysis about how humans actually talk. And a lot of it could actually be applied to computer conversations. And so, our hope really in getting together was to take some of that very academic findings and pull them into a way that practitioners could really benefit and understand them and then actually apply them in their day-to-day jobs.

Elizabeth Stokoe: Right.

[0:07:45]

Saul Albert: So, Cathy, how do you think AI has affected your history and your discipline of conversation design to date, and as kind of second question is how maybe your role has changed in conversation design especially now with LLMs being so pervasive?

Cathy Pearl: I know in my community in conversation design there was some fear happening when large language models started to really come about. I mean there had been the beginnings of large language models for a long time with neural networks and everything. But now we finally have the resources, the data, the power to actually use them, and there was a lot of, “Oh no, conversation designers, we're done for”. But the reality has been different than that. With generative AI, like I was saying earlier, I think some people thought, “Oh well, we'll just train it for a little bit and boom. We won't have to do anything else”. So, conversation designers, what are they going to be for, because all we do is write the strings, some people unfortunately think of us as.

But what we found instead is that first of all, large language models are often trained on written data, written text and not on spoken dialogue. And there are differences there. And what we do now is rather than write the exact things with the large language models say, because they generate that, we train it in terms of the architecture of the conversation. How should it start? How should it encourage the user to continue or let them know it can keep going? You know, all those things that it doesn't know because it's just a bunch of data. And so, it turns out there still is a really important role for conversation designers in shaping that. It's shifted a bit, but it's still very crucial.

Elizabeth Stokoe: And Cathy, what would you say, this is such a blunt question, but what's good and what's bad about these new large language models? And what's somewhere in-between?

Cathy Pearl: It's an excellent question. I think about this a lot. I think what's good is that, my belief is that large language models can help us be less rigid with the language we use. So, if you use a virtual assistant, you often have to really memorise, like I have to say this sequence of words and do this or else it won't work or won't understand. And with large language models, you can be very, very flexible and natural in the way that you're asking a question or asking for a request, because it's going to have a much higher ability at understanding what it is you're trying to get, your intent, your actual intent.

Another thing that I think is good is that I find I use it a lot for asking questions like, "Oh well, I heard a song at the gym and I don't know what it is. Can you help me find it?" And I start talking about it was like this and asks me questions and then it's like a cool little treasure hunt.

On the not so good side, as everyone knows, large language models, they're just like language calculators. They're just text producers. They don't know or not know anything. And so, because of this, they will produce words that are wrong, they will produce facts that are wrong, they will, you know, we call it hallucinating, which is a little anthropomorphising, but essentially they're putting data on the screen that isn't necessarily accurate. And I think unfortunately, even though all the products say, "Oh don't necessarily believe what's written here," of course people do anyway, because they've got to trust computers and the knowledge they produce. And so, I think that's the potential that people just have to remember take this with a grain of salt. It's a great starting point, it's a great thing to begin your journey, but then go verify that this is really true.

Saul Albert: Cathy, we've done some work together already with the Conversation Design Institute. We've given a number of talks and I think that's been a wonderful kind of opportunity to bring some of the insights from our collaboration to that

community. I remember, I mean some of them have been about interruption, about personality, about error and how those things are handled. What do you think are the kind of key points you would take away for that conversation design community and from that collaboration with the CDI?

[0:12:02]

Cathy Pearl: Yeah, I really enjoyed those sessions. I think again what's nice to me is like when we did interruptions. So, you and Liz gave a great overview of interruption in human conversation and of course, we talked a lot about even what is interruption because some people think interruption is when I talk over you and I'm breaking maybe the social rules. And we started talking instead more about, like I learned some things like overlap. Overlap is when I'm talking and you're maybe going, “Uh-huh,” or, “Yeah,” or whatever. And so, you're technically interrupting me, but you're not. You're in fact aiding my conversation. And I loved having these new terms to talk about with my colleagues and say they're different. And we have a term called burgeon, which again we associate more with the classic interruption, and it was really helpful to have more tools in my toolbox to talk about, “Oh, we've got overlap here, which can be good. How are we going to handle it? We've got interruption here. What do we do here?” And really being able to sort of label and talk about these different pieces of interruption even which is you think such a simple concept, but has been really helpful.

Elizabeth Stokoe: Just going back to a bit about our good and bad, what's good and bad about LLMs. What critical questions do you think we should be asking when we are introducing or using or adopting AI in our work?

Cathy Pearl: You know, I think there's a lot of sort of hysteria in the media about how AI is coming for us and it's going to take over the world, and that's not something that keeps me up at night as someone who works directly in this field. Instead, I worry more about what I said earlier about making sure people understand that the information they should really double check it. But also, I think we need to spend, and people are thinking about this in the AI ethics field, but

spend more time thinking about what is the impact of this technology on people like creators? If I can go to an LLM and create an image or make a presentation, what are people who do that for a living going to be doing? So, I think we need to think about that impact. I think we need to think about, of course, inclusion. Obviously the data isn't necessarily diverse enough and the LLM may make stereotypes about people or leave people out completely. I saw an LLM, I asked it to write a story of a boy who's having an adventure, he's in a wheelchair. And it was just like, “Oh poor so and so, he was trapped in his home,” and it was like, “No, no, no, no, let's not play into that”.

So, I think those are the kinds of things that are the clear and present things that we need to worry about and not so much that it's going to take over my home and that kind of thing.

Saul Albert: So, I guess another of these scare stories that we often hear are around, I guess, broader ethical questions around AI. I guess privacy, also social justice, you've just touched on that. But I wondered what your view was around that sort of focus on the ethics of AI more broadly.

Cathy Pearl: Yeah, I mean it’s such an important topic right now and I am pleased to see that a lot of institutions have an AI ethics group and people who are really thinking about it. I think some of it is known and some of it is unknown. Now, what is it known? I think that again we need to take a step back and really think about with new technology, people get so excited and they just start saying, “We're going to just start doing stuff and we'll figure out what it's used for and what people love it for,” and don't necessarily think about like, “Okay, well, what are we impacting maybe by accident negatively?” and really stopping to think if we put this out in the world, we can think about the benefits, like, “Oh, now someone who let's say writing is very difficult for them and now we can help them create an outline or write a better email when they're applying for a job,” or something like that, you know, those benefits. But then you think about, “Okay, well, now let's think about it in schools. How are students using this? Are they using it as a tool to help them do their homework or are they using it as a tool to do their homework for them? And

so, it's really important, for example, that teachers teach their students about this technology and say, “It's here, here's how to use it the right way as a tool and not as a replacement for your work”. So, things like that.

[0:16:42]

I know there's some institutions are like, “Okay, just hands off, nobody can use it”. But really I think we need to teach people how to use it correctly, ethically and things like that. Yeah, so those are some of the things I think about.

Elizabeth Stokoe: Yeah. And those are the kind of things that we're exploring together, the three of us, in our writing. If you're interested in learning more a little bit about conversation analysis and conversational technology in this intersection that we're trying to work through ourselves, you can look at our work on Medium, you can look at it in some journal papers. You can also look at Conversation Design Institute’s YouTube channel where our expert classes are hosted. And we really hope that you've enjoyed this little introduction to our collaboration.

[End of Transcript]