

Wittgenstein and Communication Technology – A Conversation between Richard Harper and Constantine Sandis

Richard Harper, *Lancaster University*

Constantine Sandis, *University of Hertfordshire*

Abstract

This paper documents a conversation between a philosopher and a human computer interaction researcher whose research has been enormously influenced by Wittgenstein. In particular, the in vivo use of categories in the design of communications and AI technologies are discussed, and how this meaning needs to evolve to allow creative design to flourish. The paper will be of interest to anyone concerned with philosophical tools in everyday action.

CS: When Danièle Moyal-Sharrock had the idea that, for the BWS's tenth anniversary conference, we should talk about Wittgenstein in the 21st century, we thought we needed something on technology and who better than Richard Harper on information communication technology? Professor Harper has led research groups at Xerox Europarc, as well as Microsoft Research in Cambridge for many years. He founded and directed the Digital world Research Centre at the University of Surrey. He is now the co-director for the Institute for Social Futures at Lancaster University. So maybe he could talk about the 22nd century as well! He is also a partner in Social Shaping Research which he may tell you about a little later, if you ask him. You are probably all asking what this has to do with Wittgenstein (or maybe not...) and so I thought I would begin by asking Richard how he ended up here. With this kind of background, what brings you to our Wittgenstein Society? Why are you here Richard?

RH: As an undergraduate in the late Seventies, I did sociology, amongst other courses, at Manchester University and there we were introduced to Wittgenstein. His philosophy was viewed as an integral part of understanding social science. Winch in particular was our mode of introduction to his philosophy – in his *The Idea of a Social Science*. There was a main course for every social scientist which was called *Mind and Society*

taught by Professors Wes Sharrock and John Lee; some of you here will have met them. They are sociologists and were interested in Wittgenstein and Winch for two or three reasons: one to do with the possibility that one would need to be careful about the categories used to explain things in the social sciences. The danger in sociology particularly (as in all the social sciences) is to be hijacked by categories, and Wes and John taught us to be very sensitive to the need to be adroit in the use of categories – not just how you label social types but how you understand social actions, you know, the way you describe them, the categories that pertain. And they thought that particularly some of the arguments on the history of religion, which were very pertinent to sociology, and which Wittgenstein had explored in some of his work and which Winch also explored, suggest that we need to be especially careful about explanation of human conduct when the thing about that conduct is a rather precise, rather particular idea related to a bunch of ideas with particular shapes and relations to action. Going to a church is not just something done from habit; it fosters a feeling in the body that allows a willingness to touch god. That was one of the examples of the sort of careful, precise set of relations to actions and ideas in society we were warned to look out for. Wittgenstein's remarks on the mind, mindfulness and 'mind in action', were very encouraging to sociologists, we were taught, because they suggested that perhaps *mind* can be thought of as a socially organized phenomena made relevant in action, or made irrelevant in social organizations, in such a way that some things about the mind became private, some other things public in socially demonstrative ways. Essentially the phenomenon, 'the mind', was something society arranged and was manifested in language use and that is social use. Language is society in this view, part of what makes it. Hence mind could become a topic in sociology – *mind and society*.

We were also taught that Winch offered the most explosive argument of all about social explanation and theory – in his claim that descriptions and explanations of human conduct could be – are – internal to one another. So, to properly describe an action is to describe it in terms which make sense to the parties involved (putting it simply), and this really begged the issue of what sociological explanation would become since hitherto – before Winch – sociological explanation had been governed by the principle that sociologists know better than those they study, the *hoy polloi*, the geyser on the street. The people we were looking at were fools, on this traditional view, subject to ideological trappings and such. We sociologists were in ivory towers, on elegant seats and, through our judicious reading and our pre-Wittgensteinian philosophies, knew better than the *hoy polloi*. We could tell them how social structures governed their lives, what they thought, did. But then

what we found with Winch was someone saying perhaps we need to understand the *hoy polloi* from their own perspective; that understanding how they do their world would be in terms relevant to them since it was those terms which gave that world *its sense*.

Now this contagion of possibilities – this new view – was wonderfully encapsulated by an extremely difficult person and an extremely difficult sociological writer called Harold Garfinkel. Garfinkel was not a Wittgenstein reader, in fact there is no evidence to suggest he ever read Wittgenstein. But what he did have was a genius to realize, in his own research in the 1950s and 1960s, that talk seems to be integral to action and that to understand language categories in talk you needed to understand the action in question. This was emphatically what we were taught in *Mind and Society*. This is what we were taught Wittgenstein argued, though Wittgenstein was not too interested in empirical topics, or sociological ones like Garfinkel. In any case, if you want to understand the action in question, according to Garfinkel, you need to understand the talk; the two went hand-in-hand. And what that meant was that if you are going to do sociology, if you are going to do field work, say, like an anthropologist, do not just go and look at the social structure, as it were; do not look at power, gender, organization processes (leaving aside how you might do that): look at talk, look at talk in the workplace, say; and then see how those things, like power are manifest through talk. And through the talk, in the things people do.

Now that's not the only way you can do sociology. Because even after Garfinkel, it still makes sense, in sociology, to ignore talk; and you can, perfectly reasonably, treat the world as populated by agents who do not tell their stories and look at the world that way. You know, various forms of statistical approaches do not need talk to make interesting observations about social patterns – dynamics in income for example. Yes you can actually produce very interesting research that way but we were taught another view – and that is to say, No, do not do that, instead look at talk as part of the agency of those you are studying and that led me to do my PhD. I looked at chartered accountancy, at the work of auditors in a company called Arthur Anderson, gone now, of course, and what people talked about in their work. How talk was part of their work with numbers. Shall I say more about that?

CS: Yes, why not. And maybe say something about doing things with Wittgenstein rather than commenting on Wittgenstein, which you alluded to with Garfinkel. And what that means.

RH: Oh, okay, so listen. I have heard two or three really great talks today. One of the things about Peter Hacker's talk this morning, is that he showed what I noted in my PhD, and what I try to teach my PhDs

and undergraduates about social organization, about people interacting: it is that when they talk, when they interact, they do it praxiologically, yes, in time, but they also do it collaboratively, they do it procedurally with each other, they learn about how to do ‘talk things’ together. What I mean is that talk is not just part of action, as Wittgenstein would say, but talk is done collaboratively. This sounds daft because it’s so obvious, but what I mean is that language use is essentially a sociological phenomenon, not a cognitive one, or even a philosophical one. Eh, what do I mean? So, even now, I can look in wonder at the still faces in front of me – here and now – and see that those still faces, the audience’s faces, are showing a willingness to listen to me. It’s very flattering. So my point is not that, the flattering; it’s that we are doing this collaboratively, this lecture conversation thing, and that also happens when people do anything together, when they are talking, say, using language; they do it together, and if they are doing it together then one wants to go and look at how they do it together. One does not want to look at language use from one person’s point of view, you need to look at action with words as, em, joint action.

CS: Aren’t there any exceptions?

RH: So let me get back to my thread. I was thinking, earlier, why is this being presented to a group of Wittgensteinians as a possibility – you know that language is a joint enterprise? Surely they will have read Harold Garfinkel, surely they will have read ethnomethodological studies done in his name, which are all about the procedural arrangements of talk and language in action, about how the mind, for example, is made manifest in terms of talk, manifest in things people do together, in the dance of mums and babies and they way they sing to each other in their ums and arrs, in fathers and sons and how they challenge and mock, in those interactions which are praxiological though time. First this then that, in sequences of words.

In my research that has often been the key topic I have been interested in, given that’s how you should understand social action and how it’s organized; given that it has a procedural form; a collaborative affair in language; given that you are disciplined by a desire not to be tempted by irrelevant categorization; given that you are also disciplined by a desire not to allow mysterious categories to help explain things other than the categories that are intrinsic, internal to the setting; namely, talk, turns, chat, being together in words.

So, then what do you come up with? I mean, if this is your view? Now what you come up with are, often, fairly prosaic things. To run ahead, when I first went to work at Microsoft Research, they were really keen on me helping them explore why people communicate. Because

part of my career has been exploring how people communicate, writing books about how people communicate, writing a whole book, imagine, on why teenagers text, so, absolutely I was excited at that prospect, can you imagine going to a company with billions of dollars and say this is why people text? So I did a big presentation on communication, and this is, oh was, 10–15 years ago, and this was the topic. What I was trying to explain to the engineers then, in that presentation, was that there are lots of reasons why people might communicate. Sometimes it's instructive communication. When you sit next to Constantine, for example, then you listen, you are going to be instructed by him. Sometimes, for another example, when you go out with a friend, for example my best friend Phil, he will talk at me and it's not instructing, he's just talking at me and I just listen. I am a bucket for him to pour words into. He does not instruct me about anything. For me it is just a listening. But it's different from what I do with Constantine. A listening can have different purposes. And so I said to the Microsoft engineers, the *Outlook* team, 'What could the future of communication be? You could say that what you've done with Outlook is build a communications channel, which is like a vision of a cybernetic exchange, like Locke's vision of communication, telementation. Stuff is in here, in the head of the user, and the email is used to send it, that stuff, over there, to that person's head. An organizational actor over here has stuff that needs to go to an organizational actor over there'. And I said 'Okay, if you have done that, is there some stuff you have missed? Other kinds of communications doings? And I suggested to them that one of the things you might say when you look at teenagers, when you look at them talking to each other is: What do you see? 'All of you in the Outlook team', I said, 'presumably are old enough to possibly have had kids, to see kids and to wonder, when kids are standing at street corners, what it is they say to each other? And if they are not at street corners they are sitting in the living room playing games, on-line games say – what is it they say to each other?' I suggested to Microsoft Outlook team they just pause and think 'What is it that teenage boys say to each other?' And I paused, and I said, 'Nothing, they are saying absolutely nothing'. Now the moral of that is not to be contemptuous of them and say what idiots we have produced! I was saying to the Outlook folks perhaps you might think of a communication technology that might allow nothing to be said. After all, that's what these teenagers do, so why not support it?

So what on earth would that be?

Take the possibility that when you send a message, the whole purpose of the message is this (Richard slaps Constantine on his back gently): it says 'Wake up, come on, look at me'. And then a reciprocation – Constantine does something back. Now the key property of that mess-

age – the one I have just sent – is not that it is content free, but something else: another property of that message is that once it is sent it does not matter anymore; you do not need an audit trail of slaps on the back saying notice me, notice me. Likewise with teenage boy's talk – what you need is a vehicle for them to prod each other and prod each other back, just like I prodded Constantine. What you want to do is get a little economy of prodding which is what teenage boys do and if they do not punch each other, as I just did with Constantine, they hit each other with words. The words have no meaning except the dance which is 'notice me', 'notice me'. That's the category that conveys, describes that type of communication.

So I suggested to the Outlook folks that they might engineer a communications application for a mobile phone called *ephemeral messaging*. When you send a message, the message can be 'timed out' so that it would disappear after, oh let's say, 30 seconds. It might have other properties like when it arrives on your mobile phone, this is the one that really upset them, the sender could choose the ring tone that announces its arrival. So instead of the receiver deciding how their phone rings, the sender chooses how it does and sends that ring tone as part of the message – so when I send one to Constantine, it would be 'Constantine, get off that book, get off that book, get off that book! Stop reading!' or 'Constantine, wake up! wake up!' You could say really offensive things and I suggested that part of the fun of being able to do that might be that teenagers, for example, and youngsters more generally, would know they could use this communications thing to be offensive. But not really offensive, because it's like the currency of their talk, it's what's playful to them and they know what is an appropriate thing to do. So they might just send a message for the silly sound the ring tone makes like 'whoop-whoop-whoop' because it will make them, the person they send it to, snigger and giggle. It might embarrass the recipient because Mum and Dad are around, say. My point is that this might be a property of communication pattern, a dance, a turn talking when the turns did not do anything except make play. A turn at play. And I explained that this was pointless, in a sense, this kind of messaging. But that this was its character, how it needs to be understood, its value, what it was, the category of forms of communication relevant here.

I explained that engineering such an application might be quite a complex thing to do because you would have to make it so that your phone would have control over a remote phone. This is not what is normally inside the communications application, so here you would need to do something that is normally external to the communication application, namely, control the auditory function of the phone, and that kind of thing is normally handled in the operating system of the phone, of

the remote phone, the one you are as it were calling with this messaging app. And the engineers laughed at me and they said ‘Why would we want to do this? Surely we would want to teach teenagers to do Outlook, to document their communication through purposeful communication. Why would you want them to mess about?’

And at that moment I wilted, and I thought, I’m not going anywhere at Microsoft. My attempt to convey what people do with communication with a subtle use of categories – where one use is to do playful, pointless messaging as a means of noticing each other and a category that is thus different from what you might call the category of functional messaging, the sort of thing done with an Outlook email – had not been understood. Or if it had, had not persuaded the engineers. This was not a category they wanted to honour with code. Of course, maybe my example was not right, I should not have offered such an extreme contrast. But I thought they might see how teenagers do not do functional communication in their own families and so this might make sense to them. But it did not.

That was my first big presentation. Five years later, *Instagram* messages are the common place that acknowledges this very category of being in touch. Be that as it may, what I’m trying to say here, is what I learned, and the discipline I was taught ‘by Wittgenstein’ at Manchester, was to have the good grace to try and see what people are doing when they communicate, try and be in touch, and to learn what the perspicuous description of it would be. To understand, to look at boys with each other, and think, actually if you listen to them as a parent, what you see is just *boys being together* and it does not matter if they are speaking Latin or whispering hoots or grunting like monkeys when they are there, what they are doing is *being together*. And then you can create opportunities – technological ones – for that form of being together – something like a playful, ephemeral messaging platform that would let them feel as if they are together even when they are apart.

CS: You basically described poking before, which later become a thing on Facebook...so I used to have a special ringtone for when my mother rang but then she rang me accidentally when she was next to me and that wasn’t a good idea, so it can backfire.

Audience: Are you going to share that with us?

CS: I’m not going to tell you what it said. Okay, we’ve had discussions earlier today about the empirical versus the conceptual and describing conceptual truths and making empirical observations. Do you have views on how these two come together in your line of work?

RH: When I was an undergraduate, like a lot of undergraduates, I was not really sure what I wanted to do. I was encouraged to do social

science and social science sort of appealed, a bit more than philosophy because it seemed a bit more applied, but it really was not that applied so when I had the chance to go to corporate life, to make, build, to engage with making things that seemed appealing. So, when I had the chance, first of all, the first thing I did was phone my Dad because my Dad, who was an engineer, would say to me and indeed he did, 'Oh, you will have a proper job then?' He meant, though he did not say it, a proper job rather than being an academic and sociologist who just wrote about things. Also, he was enchanted by the possibility of making things.

Now what's this got to do with this question? Oh, what have I learnt, how the empirical and the conceptual come together? Well, what I learned in the business of spending 10–20 years in corporate life, when you are working with engineering, what happens is you come to learn that you can be really cack-handed with your account of the phenomena and your understanding of language – despite your training. And when I say cack-handed I mean to say the following: one of the things you have to worry about when you are working with a team of engineers is what motivates them. Now when I did my big presentation which failed to persuade the engineers that they should build ephemeral messaging, I did not really take seriously that if you get up on a Monday morning and you are coming in to spend 5 days doing some software engineering, then you need to have something that makes you think this is 'not a really miserable job' and what motivates you is the aspiration that you are doing something good and powerful. If you say to a bunch of engineers who are earning hundreds of thousands of dollars a year, 'Hey guys why do not you come in and make an application that teenage boys can use and it's going to have no residue, no social benefit, it's just going to let them play endlessly?' they are going to look at you and think 'Oh yeah, great. Go and get a small start-up company to do that, but I'm not doing that, I'm going to do something important'. And it's not that they are pompous, it's that they are genuine and they are decent people who want to make a difference. One of the things you have to think about is what are the kinds of motivations they have. They do not admire teenage boys and their giggling talk, they want boys to grow up. So that colours their motivations, the engineer's motivations.

Let me give another example. One of the problems in the explorations of concepts and categories I have encountered has related to projects looking at memory and how you can replicate memory. Now the engineers in computer science who I have worked with on this, have imagined that human memory is like a machine memory. Now, you might think of all the classic arguments of Wittgenstein, including arguments made famous by the homunculus argument, and the disdain that that philosopher has for the idea of memory as container, as a thing that is

filled, which is pretty much what this computer science vision of memory is. But, okay, when you are working with engineers who are motivated to build something, if they have a vision that they can make a computational memory like a human memory, then justifying the thing they can build on their computer by saying it is like the 'computer mind' is what gets them going. It gets something built. And what do you end up with? You get a big, visual data processing database. That's the kind of thing they build, a data base with images. So, that seems a like a long way from the concept of mind, sure, but it is something. It is a functioning machine. And it is also incredibly fast, it's incredibly well indexed, you've got all kinds of relationships between these visual objects and ways you can string them together, and so it's fantastic in lots of ways. In your view, it's not 'the mind' or whatever that might be, not a thing, but what it is, is a doing, yes? It is a really interesting machine they have built that *does interesting things*. They would not have built it if they did not think that what they were building was just like the human mind.

So what I have learned is if I had said to them beforehand: 'No, this is not how the mind works, you should not use that metaphor; you should think of the memory as a verb, as a place you visit, as a place that you assemble, a place that you make relevant in chit-chat and talk. And, further, that you judge people's powers of social competence by the way they invoke the past not by their objective truth with which they go to the past', they will say (the engineers): 'I do not know what you are talking about, Richard'.

My point is that the engineers have concepts that let them act and those concepts that let them act are, I think, the *right concepts* for the work they are doing – the *work of engineering*. The problem for people like me is saying 'okay, when is the moment when those right concepts are no longer right for the job at hand, which is the moment when they have built something and when we need to start thinking about interfaces for it'. So you then ask 'How might you use these sorts of data stores?'

If you use such a system, let's take a thing called *Sensecam* – a wearable automatic camera device that could take thousands of pictures at various frequencies during the day like a visual memory prosthesis. A standard conception related to this device, one related version, is that you just need to remember a moment, say, when I first met Constantine this morning, and this device, this Sensecam, would have indexed that little moment and you could retrieve images for that moment. We suggested, and tried to persuade the engineers, that that was not quite the right way to treat Sensecam. It begged questions about what indexing would mean here. Surely a person would need to know what it is they are looking for if they were to find it on Sensecam, so, it was like, we

are not sure Sensecam can do anything for you except be a resource. A resource for what? It is not a memory. It's a data store. So what you needed to do was allow that database to be something that you might visit and use as a kind of resource to bring 'the past' – pictures of the past – into talk; or, let's say, when you are sending messages. You marshal it to make images available when you need them. Or, perhaps, for another possibility, you could use Sensecam data to construct little 'days in your life', a series of pictures brought together like a cartoon, or that you might send to your partner as artful renderings, as a kind of poetry made of visual lyrics, as it were, of your circumstances – a writing in pictures. And we proposed that you would have an interface for that creativeness and we would need to devise that.

The point is that there would be a moment when you wanted the engineers to move towards thinking about these purposes and novel interface ideas; that it would be *at that point* that we needed to start persuading them that perhaps, then, at that moment, their psychologistic kind of Turin Theoretic notion of memory was no longer applicable.

So, what I am saying is that I have learned over the years that it is best to be careful about the applicability of categories and concepts. One needs to *let them do* various kinds of work. Rather than saying this concept is wrong here, one needs to accept that some concepts can be right in all sorts of strange and interesting places if you understand what right means there and then.

And I hate to say that one of the things I have noticed over the years is that those trained in philosophy can have a tendency to elevate their powers of control over categories, as if they have a kind of imperial role across everyday life so that they adjudicate the meaning of categories. In my view, this is not how Wittgenstein tried to encourage us to see categories as diverse, and to be understood in use, but goes against it somehow – it seems to imply that categories should be systematic and tidy, imperial, and philosophers can adjudicate about this role.

Well, actually, lots of categories, when they get tied down to moments of action, to different places, do really particular things. From afar, philosophers can look at categories and complain that they should not be used in some way, but they do not know how that works, that use of a category, in that place; that if you understand them in action, in that place, you can see the credibility some uses might have. This does not mean you do not also need to show some finesse about when their moments of change might occur, when somethings being built on a basis of a motivation related to one category turns into something else which you then need to redesign around a different category.

But my footing, as it were, in Wittgenstein, has given me the courage to say, *allow the action to speak*, that's what you should be thinking about,

what are you trying to do next? – that’s what you should ask. And once you can be clear about that, then you can appropriate the right set of categories to help you get there – to that next step. And this can mean you lose some battles on the way.

CS: I remember when I first came to visit you at Microsoft, thinking why on earth is this guy interested in me, and I walked into your office and saw Malcolm’s Memory and Mind and some of Peter Hacker’s books and then the penny dropped, but I want to ask you, did you convince Microsoft? So this kind of argument. . . look if you think that communication is all about exchanging information for functional purposes then you are going to sort of limit the kind of software or devices you build, and teenagers – or whomever – might want to communicate for all sorts of reasons that don’t fit the Lockean model. Are they kind of getting convinced of this sort of thing now? You mentioned Instagram.

RH: Well, not always. I have got patents for key aspects of what Instagram does. So that is good, but no, we did not make that product. In the kind of research life I have had, you get used to – as do those of you who are sturdy academics – you get to know that it is not necessarily whether your arguments are right or whether the paper you have written is perfect, it is not necessarily always whether it gets published that matters, but you stick to your arguments, you keep putting them forward, you do not give up. You do things for faith, in faith that it is in the right direction even if it does not always get noticed.

I have singularly failed to influence Microsoft on its communications models despite the fact that I go to a Head Office of a company which makes billions and offer them what I think are ways forward, ways that I think my research shows are opportunities. One of my best moments was when I went to the head office of Outlook, the engineering team, and in their brainstorming room, what they call (they had a fantastic Californian name for it) *Skunk Works Room*, that was it, they had pictures of a magazine I had written, all over the walls – a magazine on why people communicate. And every example of why they communicate taken from that magazine on their wall – pinned up there – was not about functional exchange of information; it was about every other thing you might do! And they had it up there as their blueprints for the future! Including my ephemeral messaging concepts! None of those ideas ever showed up on a product. But I was gratified they were thinking about them; well I assumed they were. My ideas got to the reviewing table.

But there are other examples where I do think I might have influenced technologies more successfully. Some of which everyone in this room will have used, I would think. I can think of something that I helped evolve but in some of these cases I do not have any credit for. I am not

complaining about this, this lack of credit, because good ideas – and I am not saying mine are always good ideas – but rather, many of the ideas I toy with are common currency. Good ideas are often common ideas. So, in the case I am thinking of now, we were trying to persuade Bing that the problem with the search engine user experience was that it is conceived on the basis of the human user as merely an information seeker, like a computer trying to find a piece of target data, and so the systems are optimized around getting an answer to a question, data for the question, as quickly as possible and the experience of using a search engine reflects that, the user interface and such.

Indeed, that model was used for search engines at first; indeed one might say that was actually what those early users were doing. It really was a mixture of FBI agents and engineers in Palo Alto trying to find key pieces of information as quickly as possible. So, the model of the information forager was right at that time. But we tried to persuade Bing that if you looked at people searching today, years after the first search engines were built, say when search engines are being used now, today, on a Saturday morning for example, at home, what you would find is that the users seem to be doing something else. Not the same thing as these early models assumed. People seem to be using a search engine to *pass the time of day*. So, as a measure of this, there does not seem to be much logical connection between one search and another search.

When we talked to people who were using search engines on Saturday mornings, one of the things that they said to us about this strange voyaging, is they were just *playing on a search engine* and having fantasies about holidays, fantasies about sports cars or, at my age, a fantasy about motorbikes I suppose. And then they said ‘I wish you could keep that somehow and make you know, links.’ So we came up with the idea that when you did a search you would not just have the kind of URLs or SERPS (Search Enquiry Results Page) as they are called, but you might do something else, which is create what we called a *Card*.

Some of you will be old enough to know what a cigarette card is; some of you, if you are North American, will know what a baseball card is. It is like Constantine’s notebook with a picture at the top and some text. So, we came up with the idea that when you do a search, the search engine would find websites that seem most likely to provide what you are looking for, but also go through those websites, take what seems to be the most likely picture and the most likely text from them for your search query, and create a *card* and present that card on the right of the search, like a little Wikipedia entry for that particular search, as the answer to your search. And we suggested that you could consolidate that – those cards – bring them together and keep them and people could drag them onto their

desktop (from their browser) or they could send them as an attachment in an email.

So, to put this another way, a user would do a search and make a little content, a little card in their search, so it would be a picture of, say, Peter, and some text from a website that has been scraped about Peter, and the card they had produced would be unique, an output of their search. And if they quite like that card, they could say 'Look, I want to send it to a friend'.

And we said we can do this! I mean, we said to ourselves we could make this application.

So, we did a big presentation in San Francisco to the Bing team and before we went there we got lots of patents and working prototypes and things to make our ideas seem real and profound. Anyway, we did the presentation and there was absolutely total silence, it was very weird. I mean, absolute total silence. No one said anything. And we were rushed out of the meeting room 10 minutes after our slot. And this presentation was to the head executive of Bing who is now the director of Microsoft research. Okay, we thought, they did not like the idea, he did not like the idea, that is why the silence. But 3 months later Google came up with the same thing, as did Bing; exactly the same thing! *Cards!* So we came to think that maybe they had been thinking just as we had at just the same time. The silence was actually an affirmation of our idea, but they did not want to tell us.

At that time, we did not know that. That our idea was on the mark. One of the reasons we got silence, we thought, we thought at the time, is that when we proposed these cards, some of these cards would not select the best picture of Peter [Hacker] or the best text of his or Constantine's works. There would be errors. And we assumed that the search engine people do not do errors, they are in the business of 100% perfection and so our error rates were not acceptable to them. And that was the reason for the silence, that was what we thought, since we were proposing something that was not on their path to the Holy Grail – perfect content, perfect answers to search queries. That is how we walked away from that presentation. But, obviously, we were wrong. The search people do not mind errors, if they think the user does not mind, or is doing something else, if the *category of doings* is not like the information forager model but is something else, like journeying or gathering, card collecting, in this case.

So, today, if you put in my name in any search engine you will get a slightly random picture of me and some text scrapings – lots of good examples one can find – but anyway there will be a card on the right of your SERPs page. You might not know it is called a card, but that is what it gets called. That card is slightly lashed together, it is not perfect and might have some wrong text or the wrong picture, the wrong

Richard, but people like cards. Mostly, they are good enough, they do the trick. Millions of them are produced every day!

So, you know, I have had success and the trick about this was to try and suggest to engineers you might make them think differently about the uses and the grammar of action, the grammars related to categories of doings, searchings say, time-usings, gatherings, things you can do with search. Is this what you wanted to hear?

CS: It's what everybody wanted to hear. Just before we open up to questions we have a few more minutes I wanted to hear some things about meaning and understanding as well. Possibly in relation to AI.

RH: My last big job, just when I met with Constantine, was working with the speech engine folks at Microsoft, working on smart systems like *Cortana* and things like that. And this is illustrative of a number of things I have already mentioned. The speech engineers were completely enchanted by the possibility that you could build a speech engine which is 95% more accurate in terms of reasonable interpretation to an utterance of a user. And on that basis, they felt like saying we could change how people use machines and their aspiration was to maximize the accuracy of utterance recognition and, on that basis, maximize the sort of response that machines provided. Having done that, they said, surely a person would then be understood by the machine and then a person would be able to understand the machine, so then you could have a *conversation between a person and a machine*.

And we said to them no, no, no, no, we understand the motivation you have got, which is to be able to talk to machines; that is a simple motivation, a kind of categorical criterion that you can use to say how can we pass this. People can use any utterance, more or less, and the machine will come up with some sensible response. That is not the same as saying, it is a different category from, people *having a conversation with a machine*.

Firstly, we said, or rather we tried to persuade them, what would a machine say to a user that would make for conversation? If I had asked the machine to find Constantine, and where he is, if the machine were to say 'Here is Constantine', would one call that a conversation? Well, it is certainly a verbal command and respond cycle, but a *conversation*? Is not conversation about something else, interiority, perhaps, sharings about inner thoughts? Is not that what one might say is a conversation?

In that case, if a machine entered into a conversation, and sticking with the same example, would it say, conversationally: 'I used this probabilistic technique, combined with this proportion of your prior search query cache, combined with this geographical location, to identify the kind of information you want – namely, where Constantine is.' Is that the *interiority* of the machine that would make for conversation?

When we doubted that this would make for a very good conversation, the engineers got very vexed and said no, no! So, we responded and asked them to take seriously what would conversation with a machine be about? What is it you are trying to achieve when you are having a conversation? We asked whether they were misappropriating the term conversation when really all you have, what you have got with things like Cortana, is a speech dialogue system. An input-output procedure. And if the speech dialogue system works, surely you would need to understand why people want to speak to a machine, what benefits them in doing so, why they prefer sometimes to speak to a machine and sometimes prefer that a machine speak back, so to say. What the input output is for. We asked, what would be the situations of use in which that desire would be applicable? That dialogue? If they elevated their aspirations into something which sounds like the question of understanding, which is often what is being implied when the term conversation is used (its grammar if you like), we said that it becomes a different sort of thing.

What we did try and persuade them to accept, and Constantine and I are working on a paper on this, was to look at how *understanding makes sense* in certain places, is a sensible purpose, whereas in other places, or doings if you like, understanding has a more profound sense. It is also sometimes a word that labels something that can be subtle. It is one of those wonderful categories where you can use it as a label to point towards the evanescent – somehow the opposite of understanding. So, for example, if I were to speak to my Hoover, I would hope it would understand what I mean when I say ‘Can you Hoover my carpet?’ When I ask one of my teenage boys ‘Do you understand what it feels like to be a vexed father?’, I do not expect him to say yes or no. I am saying something about the relationship, I am using the word ‘understand’ differently. But if that example is saying anything for the engineers using speech systems, it is saying they have a need to make a taxonomy of those tasks where the understanding which is produced between an interaction and a user and a machine is the kind of interaction where the understanding that is required is achievable by the machine and desirable on the part of the user. Does the user, do I, want to find that my Hoover is like Oliver, my middle boy, who yesterday morning was so hungover that he did not know what house he was in? I do not want my machines to be giving me unnecessary interiority, I do not want to understand in that sense. I genuinely do not want to know my Hoover is hungover.

So you need to come up with a set of criteria for distinguishing where the meaning of understanding is intelligible and applicable and those places where understanding is one of those phrases which label the fantastic, awesome, wonderful richness of human endeavour. When I say

things like 'Do I want to understand what my wife was thinking when I came down here on a Sunday morning to a Wittgenstein conference?' Probably not. My point is some words, some categories, are very powerful, very useful but you need to be careful with them and in my business, my career, navigating doings with words, conceptual analysis with words, and trying to get things done is, has been like a handicap race which I have fallen over lots of times but it has been a great pleasure to compete in. Understanding is one of the words I have tangled with, as is communication and the phrase 'being in touch'.

CS: Thanks so much Richard, that's really wonderful. We have time for some questions if anyone has anything they want to ask.

Q: Have you had any experience with engineers trying to produce intelligent machines, or anything like intelligent machines like intelligent cars and the problems that might pose?

RH: Yes, though one problem here is that there are many different types of what an intelligent machine could be, and many different forms of intelligence or intelligent roles. So one of the things you might say about intelligent cars is that what the car is doing is behaving like a train and one of the things you might say about a train is that it is a wrought iron solution to intelligence because the tracks dictate where it goes. Now the problem with the car is that it is evolved in a situation where there is choice. We are now moving away from a situation where there is choice, to a situation where there is less choice. Traffic volume means that there are fewer choices to make. We have to drive in train-like caravans. So, when it comes to intelligent cars, it seems to me better to think of them like trains. So in my mind, intelligent cars are not very intelligent. That they use a lot of AI to make probabilistic technique judgements about proximities, trajectories and routes, yes, okay but to me that is not hugely interesting. It is a mere engineering problem.

There are lots of other ways to use intelligent systems where the meaning seems more appropriate, closer to the everyday meaning of intelligence. Let me give you an example of the kind of complexities though. For many years, search engineers thought that the problem their technologies solved would be, would come to be, solved by AI. That is, they would come to use AI to predict what the users wanted and to have targets available to respond to those wants. A user wants to search for this, and the search engine delivers it. The answer. Yet when mobile phones took off and, in particular, when mobile phones were internet enabled and searches on mobile phones started, it made a complete mockery of the role of AI in search because it turns out that where you are on your mobile phone is massively determinate of what you are searching for. So, if I am looking for a Katja [Behrens], say, where I am

determines which Katja I am interested in. The one I am after is almost certainly the one who lives nearby. Not one who lives in Alaska. Massively. My cache, my historical record of use, the history of my prior searches, other determinations of what I might be looking for, all of which might use probabilistic AI techniques to come up with an answer to which Katja I am after, might be relevant but nearly always not as relevant as where I am. So, the hope that AI would help search engines better determine what answers I need is not as great as we thought. It turns out where you are is more important. And that does not sound like information that requires intelligence. It is just a geographic coordinate.

So how our AI unfolds is much more complicated than would seem. As it happens, this Wednesday I have a meeting with Microsoft operating company that might have a big problem with AI. That is, they have clients who imagine that AI is going to solve lots of problems and they have machine learners from the laboratories who say they have solved everything and the two do not meet in any regard whatsoever because one of the problems the clients have is that they imagine you can make data of things – where there is no sensible way of making the phenomenon in question into the data types AI systems require. I will give an example. Some large retail companies want to know if you can use vision techniques to identify from a person's face what they want when they walk into a shop. You get machine learners that say we can read moods, so they put those two together. The big companies and the machine learners must be able to do this and, in the middle, you have people trying to make a business out of it and they asked: 'Richard can you come along and help us disentangle this terrible mess?' Because there is no way this is going to be solved and part of the problem here is that expectations of people in AI are grotesque and have become hyperbolic and often turn around really detailed possibilities, and detailed possibilities which need to be thought through properly. A nice example is whether visual processing machines can identify intention from the shapes of faces and the movements we use in everyday language expressions. But in a shop, all the camera sees when it captures a signal of someone standing still is not a person, but a still creature, it is not a lingering creature; to linger is an implication about purpose. Lingering is incarnate in action, in how we understand people but it is not visible in what a camera sees. But this is what the shops are interested in – how do you recognize that some person's movement is a lingering? If so, why are they lingering, what is it saying about them? Can you recognize that the creature has stopped moving because they have seen something on the shelves? Or because a purchase has been brought to mind, to their mind, by the placement of goods in the entrance to the shop and

so they pause, they linger? And that is a big achievement for vision systems because, when the creature stops moving, it is not necessarily visible what they are about. Visible here means two things – as in things seen and things seen as actions in the Winchian sense, as purposeful, as understood in terms of a concept of action. A vision system cannot distil the difference between a lingering which is pointless and one that is purposive – a pause to prepare for a purchase.

Q: Ok, I agree that there's a lot of fun in this area, there's a lot of natural and good fun in this area, but I cannot deny the fact that I am very troubled that this is sounding as if it's an ethics free zone. Technology is reaching a point where there's virtually no limit to what we can create to meet a demand or a requirement or wish or a desire of a market who will pay money for it. That leads to the question – and it seems to me to be a very genuine, real ethical and cultural question – where there are no such limits, how do we decide, should we simply make anything, everything because we can? Should we make anything and everything because someone we've seen is willing to buy it? This seems to me to be a serious question which is having a cultural impact. Many of the things you have mentioned, I mean my first reaction when you said about the fact of what had happened, was at last I have found out whose fault it was.

RH: I wholly agree with that. What one learns from reading Wittgenstein, is how delicate and diverse are the issues that are articulated through different forms of expression and likewise with technology: it is not to be thought of in terms of ethics in a *carte blanche* way. There are all sorts of different sets of possibilities and issues and different sets of ways of making ethical choices which are effectively bound up with the particularities of some kind of engineered, technologically enabled possibility.

So, with that as background, let me answer your question this way. When I first made this presentation about being able to control the remote phone I said to them, the engineers, I can show you how interruption in talk is a recognizable phenomenon. And I can show you as well that people respond to an interruption by noting it is an interruption and deal with it accordingly. They manage that unethical conduct, if you like. Further, I can also show you that sometimes people will allow interruptions because some things deserve to be interrupted – sometimes interruptions are ok. Now, by the same token, given this, I suggested that an ethics of interaction would emerge with a communication channel in which one could manipulate my recipient's messaging device in a way that seemed unethical, through interruption say. So, I would sometimes wind up a friend by calling at inappropriate times, by interrupting them, but I would, over time, come to settle into a pattern where most often the regime of behaviour would be subject to propriety

– my friend would allow a little interruption, I would create some. It would not just be me, either, sometimes my friend would impose controls on me, like deleting my messaging rights if I interrupted too much and this would affect me. I would have to renegotiate comms. rights. We would eventually come to a settled pattern if you like, an ethical way of handling the borders of communications ethics together. Of course, interruptions seem a long way from ethics, from, oh let us say, trolleology and the kinds of examples – risible in my view – used in that form of thinking about ethics. My point is that there are ways in which propriety and ethics can show themselves, in all sorts of subtle ways, and if you wanted to tame the big technology companies the way to do that is not to approach them in the general, not through trolleology reasoning, but approach them on specific sets of issues and specific solutions to do with, say communications, and the way forward is to show how ethical operations could be done and managed by the users: and these operations could be an outcome of good design, design that allows users to act with propriety. Managing interruptions depends on allowing interruptions, technologies of communication might allow that.

Q: Just on a similar kind of vein, one of the things that technology and AI try to do is a lot of machine learning, trying to predict what I want, and I wondered whether that was a value you thought would shift. Will there be other values that's will be more important later? It's actually a problem for me that it does that. Not only are the predictions sometime quite incompetent, sometimes they are helpful but sometimes they're not. But I don't want it to, especially Google, I don't want it to guess at what I'm looking for because that distorts my process so I then have to take quite a lot of trouble to make sure it doesn't do that.

RH: No, exactly and here's a good example of ethics, which is I think that ethical considerations sometimes need to be cut open and made a practical matter; something that a user, in my language, can act on and manage if you like, so that people can make different choices around what the technology can do. So, for example, sometimes I do not mind when the search engine uses my history and makes some predictions that can make my search quicker. But, you know, sometimes other considerations come into play. Let me illustrate. We had a big project with Bing on ethics and search. It turns out that on Saturday morning there is very interesting user behaviour phenomenon that shows itself – because search engines are often emptied on Saturday mornings; the cache is removed, people reboot search engines, well not just the search engines, but the browsers too. And Bing was convinced this was because people were doing porn on Saturday mornings and trying to hide their activities. That is why they deleted their search history. (That is what I mean by emptying the search engine. It is putting it simply but it is pretty much what

happens). And we did studies which showed virtually no one was doing porn on Saturday mornings; they were shopping. And one of the things that they were doing was getting fed up with the predictive aspects of the systems leading them back to the same set of things to buy – boots, boots, boots! But we also found that sometimes they were genuinely wanting to hide what they were doing, their shopping; they wanted to hide it from their partner but not because it was porn. But because they were trying to buy a present for their partner!

So, are there ethics here? Yes. But they are prosaic, in a sense. It is not big ethics here, but the ethics of buying presents and not being found out, and, related to that, there is the practical desire to get the search engine to work for me, and not constrain me by overly predicting what I am doing. When I shop, when anyone shops, one shops for this and then gets bored. Then, one looks for something else. It is hard to predict when that happens. And so users end up having to show the search engine this by emptying its cache, so that a new session starts.

So, is there an ethical muddle? Who is doing what here? The technologists, the search folks or the user? There are lots of ways of figuring this out and moving forward, allowing a better dialogue between the user and the system, and so forth, allowing what I just called better management. I think the user should be heavily involved in this.

But there are also these big arguments about, for example, the ethics of keeping personal data. This seems to me a different set of questions. Now one of the things you might say about the ethics of personal data is that it is a great myth on the part of the search engine companies that they do anything with much of the data they collect. They genuinely cannot do anything with most of it because – well, what do you do with it? Aggregations of these data most often come up with pretty bland implications. Someone who has just looked at a car is probably just about to look at a car again. That is what you find. You do not need rocket science for that, you do not need state of the art AI to do that. Nor do you need cloud farms with billions of instances of the same; so it is bizarre that the search engine companies and browser companies seek to gather all these user data. It is just digital dirt.

Q: Well this is possibly a question about those big data farms where they, as you point out, there is very little practical stuff that is got out of it and so going back to your question about concepts being right and wrong and your suggestion was that concepts that motivate engineers are just sort of fine so they just keep them turning over and producing this stuff and then at the end of the process we come along and adapt it and make it useful to us by giving it, applying a new category. But then there is this issue that, well maybe, the concepts that are motivating engineers have a kind of logic and destination and perhaps a religion of their own that are all heading in a certain direction and, as

it were, trying to humanize the project at the end of the process might be, might not be a failsafe project.

RH: Yes, I think that can be; but I also think you are doing a kind of sociology of reasoning.

Q: Well you're a sociologist!

RH: Now hold on, let me finish. There is a tendency in sociological treatments of 'reasoning' that exaggerates the structure and scope of reasons. It is a kind of craving, to put it in Wittgensteinian terms, this bigging-up reasons. For example, sociologists are particularly keen to give a kind of eschatological flavour to reasons; as if the reasons to do this here and now are related to the end of the world, as if reasons, ideas are all a schema. Think of Jameson, a cultural theorist but basically a sociologist. But in contrast, when I have worked with engineers, very few link the ideas that motivate some product they are building with bigger, longer term goals. Their reasons are, well, more piecemeal. You know, one of the things you might say about lots of scientists is that they have a philosophy of science and they are very keen to explore and understand what science might be and what it leads to – its big picture outcomes. The way these outcomes are described varies in different regimes of science, but the general tenor is common – think of the claims physics makes about the reality of the universe. You know the kind of scientism one comes across. But it is not likewise with most engineers. They have a very rough and ready view of reasons, and the categorizations to do with those reasons. Occasionally, you will meet someone with a religious like edifice wrapping up their work – I want to make the world a better place with the digital, you know that sort of thing – but they are rare. Mostly what motivates them is, really, things like it is a Monday or it is a Friday, and it's crude and it's simple. It is almost inane in its symbolism. But they abide by these reasons, however inane they might seem to be. I mean, their reasons are basically practical. Really, once they have done it, made whatever it is they are making, they have no reasons; they do not seem to be (I am saying this as a kind of anthropologist) wed to the reasons that persuaded them to make that thing. They are perfectly happy to be sold another set of accounts for that thing that makes sense for them, that would get them out of bed on a Monday morning. They would not take any old view, though, I am not saying that. The issue here is really to do with motivation, what provides it, ensures it. Software engineering is a really miserable job. It is incredibly difficult and you can spend weeks on an algorithm and if you are not totally convinced that algorithm is going to make a difference you are not going to

do it. It really is, it is not a simple problem, so having reasons for working on it is part of how you get work done.

But what I am saying is that this view, this view is always also somehow deeply practical. It is not pious, it is not part of some big system of ideas. I gave you an example of memory, how a notion of the mind having memory like a computer, was part of the motivational fabric that produced that wearable camera devices I spoke about, Sensecam. But the engineers did not think they were making claims about cognitive science, or reflecting some unified idea of what the mind was. All they thought was 'That is a good idea. I could build that'. And once they had done it, they were very happy to let other people come along and say, let us forget that memory thing and do this with Sensecam. But they needed that idea to enchant them, to make them want to do it when doing it is, well, hard, a daily grind. Coding is work, after all, it is not art. So, you need vehicles to motivate you. Concepts. There you go. I would like to finish now. I have run out of words.

University of Lancaster
Lancaster
UK
r.harper@lancaster.ac.uk