5._Collective_Intelligence_-_Closing_remarks

First of all, I don't think we spend enough time talking about what intelligence is. And there are many definitions of it, some of it empirical, some more almost philosophical. We all have a clear, sort of intuitive view of what it is. I mean, some people are cleverer than others. But exactly what you're defining in making that statement is difficult. And it's clear that we can make mistakes in that, too, and I think that plays into this judgement of the confidence of others which we take into account when we're weighing the significance that we should apply to their judgments as opposed to others'.

What we need then-- if we're interested in enhancing whatever intelligence is that is really useful empirically in the real world, we have to get a clearer view of what it is. And subjective impressions don't necessarily tell us about it.

You can tell that from what we feel about the intelligence of other species as we look at them. We can't communicate with them, so we can't judge from what they say to us about how intelligent they are. We don't know about their social backgrounds or their education to help us.

We can look at animals and very often we're inclined to say, oh that behaviour, that dog or that cat or even that ant or that whatever, doesn't that look remarkably intelligent? But if you burrow down below the surface of that, you often find that what appear to be clever actions have a very, very simple basis. That shouldn't surprise you. It's natural selection. And animals have to succeed in the world, so the simplest and in some cases dirtiest or crudest mechanisms in their brains for enabling them to achieve that will be selected for. So very simple algorithmic mechanisms can appear to be very intelligent.

I'll give you a couple of examples. David Ingle was a guy who studied frog behaviour a few years ago. He was a physiologist and behavioural psychologist. When he looked at recordings from the visual part of the brain of the frog, he found neurons that responded selectively, very precisely, when a horizontal extended line became shorter. And he found that that was the optimal stimulus for making frogs snap at what they thought was food.

It is a remarkable fact that frogs will not try and eat things unless they're moving. They respond to the movement, and it turned out that the optimal form of movement was the contraction of a horizontal line.

And Ingle said, well you know the obvious interpretation of that is, gosh, this frog is very clever. It knows that this is a potential piece of prey-- let's say a worm or something, a maggot-- and the fact that it's getting smaller means that it's receding in space. It's going away. So it could infer from the fact that this little thing is trying to go away that the little thing has spotted that it's a frog and wants out of here.

So you could imagine a complex kind of logical inference process driving behaviour. In fact it's a neuron that happens to have a receptor field that responds to things getting smaller that triggers the behaviour.

Another was Horace-- again, in frogs-- Horace Barlow in the 1950s was the first person to record from ganglion cells in the frog eye and developed an idea that was called the trigger feature hypothesis, that classes of neurons have evolved to be triggered by behaviorally significant features in the world.

And he found one class of ganglion cells in which the rate of nerve impulses was influenced by where an object was within the receptive field of the neuron in space, such that the cell gave a maximal response as it was heading towards the centre, towards the centre of the field. It was a very simple structure of input from the receptors that could explain that.

But he pointed at that from the frog's point of view that was marvelous because it was predictive. The firing pattern of the cell predicted where that object was going to be in the future. It was going to be in the middle of the receptive field of this particular neuron.

And he calculated from the optimal velocity of the cells how much movement in space that would be, and how it corresponded to the reaction time of the frog. And found that the frog could simply compute from those cells where an object was going to be in the future and snap at it based on this signal coming from the receptor field telling it about movement towards a point.

So again, things that look terribly-- gosh, isn't that frog clever. It's calculating where the object's going to be, knowing the processing time in its own brain, and getting it right. Not at all. It can be done algorithmically incredibly simply.

If you believe that IQ conventionally measured has anything to do with intelligence, there's an interesting fact. And that is that IQ correlates really extraordinarily well with reaction time. If you just get someone to press a button and a light goes on, and measure that, look at how it's distributed in the population, there's a very high degree of correlation with what you measure with IQ.

What's that telling you? It's telling you that we live in the real world, and the faster you get things done in your brain, the more chance you have of exploring different options in the logical processes before coming to a conclusion. So if for whatever reason-- I don't know, your axons are better myelinated, or you can get through circuitry a bit more efficiently-- if you can do things a bit faster, than within a given period of time you've got a chance to explore more options. So speed is of the essence.

If we think, then, about how we might improve collective intelligence-- and I'll make the argument in a second that that's limited by timing as well-- then improving the speed of reliable communication could be crucially important. And that's where digital technology can obviously help, because the limits on temporal processing are completely different from those of human brains.

Can we trust our intuitions about what intelligence is? is the next thing. I think we really need to set for this field empirical and objective measures of what we think that good judgement is. I don't think that we come with a reliable subjective knowledge of how difficult things are to do. I mean, how clever people have to be to do them. For instance, it's pretty clear if you look at a brain that seeing the world-- just seeing-- is much the most difficult thing that we do. It involves about a third of all the neurons in your cerebral cortex. All the back of your head is just filled with machinery for doing vision.

Another way you can judge it is by how difficult, how long it was before computer techniques, computer graphics techniques, and even worse, visual recognition techniques in computers-- how long it's taken for any degree of progress. And certainly visual recognition is much, much worse than in human beings. And yet we all do it effortlessly. We don't feel, any normal person doesn't feel that vision is difficult.

And yet everybody feels-- well, almost everyone-- feels, say, I don't know, solving a Rubik's cube or doing mathematics or doing that silly little logical thing that some of us got wrong-- we all feel those kinds of things are very difficult. They probably occupy very small parts of our neuronal system. And there isn't a simple relationship between how difficult we feel things are as tasks and how difficult they are in reality.

Another extension of the unreliability of our assessment of intelligence is that we tend to ascribe particular significance to individuals who have special skills-- expertise, talent. And in many cases that's good. If we have some kind of intuitive feel for, gosh, that person is very good at this or that, it can help us to identify individuals who are reliable when it comes to the transfer of knowledge and the group cooperative action. We can recognise kind of intellectual leadership in one area or another and seize on it and utilise it.

But it doesn't always work. The phrase "idiot savant" is I think useful. The fact that some individuals who clearly are not generally competent in intelligence can do certain things at a level that most individuals can't. I don't know-- eidetic memory or playing tunes on the piano from memory, that kind of thing-- which sometimes comes with forms of-- autism spectrum disorder isn't used these days, but you know what I mean. That kind of condition.

Those sorts of properties are not necessarily really difficult in terms of neuronal circuitry needed to produce them. We tend to admire them because they're unusual.

A lot of the discussion of course today was about whether and how and when groups do better than individuals. It's quite clear that in simple detection tasks-- just did you hear that noise or not? Did you see that flash of light?-- having more individuals just doing the task completely independently of each other-- handing one a piece of paper and saying yes I did, or no I didn't-- does better than individuals do just because of the statistical process called probability summation.

The chance of detecting a very weak stimulus increases with the number of detectors. It increases to the square root of number of detectors. That doesn't require any interaction between people at all. So just sharing ideas without any complex process of interaction can help in very simple judgments. And I think that's a lot of the basis of things like crowd sourcing, for instance.

Looking, let's say, at luminance variations in distant stars, which is an indicator of the presence of a planet. That's one of the crowd sourcing projects at the moment. No communication between people who do it. Everyone's sitting looking at their displays from Hubble or whatever it is, just waiting for it to happen, and then sending in a note by email when it does happen. Completely independent detection, but it improves performance compared with any one person just searching through all that material. So we can get large improvements in performance without anything magical happening between people.

The improvement over time is I guess what one would want to iron out. If you put a group of people together, how can one encourage that over time their collective judgments together on a particular topic get better and better? And that depends, I think, on the kinds of processes that Chris and many others described. One dependent on communication.

Now communication is used by all animal groups that act collectively, from ants to fish to birds or whatever, but it is very often simply based on one individual looking at the behaviour of another, seeing whether it's effective, getting feedback from that, and incorporating it in their behaviour. And that takes time, and time is valuable.

So you can look, I think, on social signalling-- say calls in monkeys to indicate the presence of potential predators or whatever-- as being a way of speeding up a process that could more slowly have been performed simply by observation and the feedback of success in others. And language allows us even to escape the immediacy of the situation to think about problems that might occur in the future and convey them to each other. So language, which we spend a lot of time talking about, is very important in speeding up the process.

We live in real time. We have to act in real time most of the time. Pace Dan Kahneman, Thinking, Fast and Slow, we don't have the luxury of being able to think slow about many of the judgments that we're making in life, judgments that require intelligent behaviour. So the luxury of an individual being able to rehearse, as it were, what might be achieved by collective interaction is important for operating in real time.

Finally, how to improve collective intelligence I guess is the ultimate agenda. Robin, of course, would tell us that we have to optimise the group size for the task-- maybe up to five for decision making in groups, perhaps 15 for debate, intense discussion around an issue, 154 consultation, and a larger group matching to Dunbar's various numbers. Although I guess that's suspicious because you can produce a number which matches any data there are, but I'm sure he's right. That's important in setting up environments which are optimum for collective intelligence.

But in all of these processes it's important to remember that any process that ensures that some individuals can influence group judgments is only as good as your ability to discriminate the expertise of individuals. So bad ideas can transmit as well if they depend on the charisma or the confidence or the bullying tactics or whatever of the people you're learning from. So, I mean, religions are a good idea. War on drugs, and policies which are driven just by the strength of conviction which transmits pretty effectively to others and can certainly influence their judgments badly.

So we have to think of constructing mechanisms for enhancing the recognition of genuine expertise if we're really going to optimise collective judgement, recognising faulty ideas as well as good ones.

The use of digital technology, the use of the internet, has been discussed a lot, and there clearly is great potential in digital devices not least because of their speed. Take, for instance, chess-playing programmes, which now-- the best of them-- can defeat leading chess masters.

But they achieve it by an entirely different approach to those that chess players are using, which tends to be based on a kind of perceptual analysis of the board as they see it-- seeing the range of possible moves and so on. Most of the software approaches taken by chess-playing programmes are exhaustive and just do it by massive number crunching-- very, very fast, something the human brain can't do. So utilising the distinctive characteristics of digital technology is important if we're going to use it optimally in aiding collective intelligence.

But that's not to say that the internet, for instance, simply as a medium for communication is necessarily any more effective than face-to-face interactions. Certainly looking at the terrible world of Twitter, it's just as influenced by the same kind of tactics that produce false attribution of competence as normal conversation is. There's bullying. There's flamboyance. There's just the volume of presence and all those things which can fool us in ordinary social intercourse. So it's not necessarily a good thing just to transfer all of our interactions to the internet thinking that that's bound to make them better.

Another is that the internet just conveys-- or even in some cases, I suspect, amplifies-- some of the distorting factors in normal communication. I think Licia mentioned that only 10% of those who contribute to Street Mapper are female, for instance. We're going to end up, if we're not careful, with more of the same gender biases in exposure to views and pressures in decision making in internet communication.

OK. Yes, we're nearly there. Yes, so there we are.

So can we learn-- can we learn--

Is knowing about how collective judgement is implemented in our brains or in psychological structures, in cognitive mechanisms, could it actually make us better utilising it? It's a ridiculous speculation at this stage, but if we knew about what goes on in the brain as a group of individuals work together on a problem, then if you know where it's done, know the structures and the networks involved in doing it, know the plastic changes in the brain that are going on in individuals as they are influenced by others, it is not inconceivable that we could design ways of improving those processes by modulating our brains in a way that can make them more effective.

Although the best way of modulating brains is through teaching people. That's what education is all about. It's about changing people's brains. So if we can learn the rules and just incorporate the rules for better use of collective decision making, then we might make much more effective societies.

And just finally, this point about confidence and competence which has threaded its way through all of our discussions today. It's not just a matter of enhancing the ability of people to listen to others, to take account of their views, but improving their ability to distinguish between confidence and competence. Look at the way that we still respond to our politicians, to our leaders in every field of life. We are clearly-- we are collectively very much more influenced by confidence and show than we are really by the quality of the advice we've given. That's all I want to say.

[APPLAUSE]

Well, thanks, Colin. I think that the good news if you ask the public in surveys who they trust most, I think university professors, medics, other groups, are far above politicians and journalists and other low grade roles. I think all that's left for me to say is really to thank you for that summing up.

Thank you all for coming, and in particular the presenters. I think we will, as Stefana says, try and make papers available within the limitations of copyright law and so on. But mainly we're trying to begin a conversation where there is more cross-pollination between different areas of theoretical advance-- and we've heard from a lot of disciplines today-- but also a lot of the practise many of you are involved in.

And I think some of the theory and practise aren't very interconnected at the moment. Yeah, that's probably where a lot will evolve, whether that is around modulating people's brains as they work together-- which is a lovely image to end on-- or on the specific formats of meetings and the various Dunbar numbers and how those can all be orchestrated.

So I think there will be some drinks outside. I'm not sure what the historical role of alcohol in collective intelligence has been, but there will be some alcohol on offer in about 30 seconds time. But please join me in thanking everyone who has spoken today.

[APPLAUSE]

And finally, thanks to Stefana for putting it all together with Mattia.