

## **4. Collective Intelligence - Extending the Collective Mind**

It's a pleasure to introduce Licia Capra. Licia is a professor in pervasive computing in the Department of Computer Science at UCL. And I'll introduce the rest afterwards.

Thank you, Stafana.

So I'd like to go back to some of the topics that we discussed in the first session today, with Paulo and [INAUDIBLE], on the role that technology plays in enabling and facilitating the collective intelligence processes. Where we have spontaneous and large groups of people coming together and sharing their knowledge, their skills, their data, in order to perform a collective task or goal. And these process, as we said, are nothing new, but what happened recently thanks to technology is that they're taking place at a scale that we probably have never seen before.

So during the first breakout session, we were asked to think of what are example of collaborative intelligence processes, are successful examples. And I was thinking when in terms of technology mediated collaborative intelligence process that I can think of, where these are the first that to me come to mind. So Wikipedia is a prime example of where we have everyone around the world who has access to a computer and an internet connection, who has the ability, now, to contribute to gathering and maintaining the whole body of knowledge.

But there are many other examples, such as a Stack Overflow, a question and answer forum, where programmers and software developers can share their knowledge, as to help each other. OpenStreetMap is yet another, very successful, example of where people look and turn themselves into cartographers, and contributed to build a map of the whole world.

And what technology has been a particularly successful in doing here is to break down the barriers that separate people in space and time. And everyone around the world, as I said, just with a computer and access to the internet can come together and contribute to these collective tasks.

And what is quite interesting to me is that the very same technology, which is enabling this global scale collaboration is being used also to perform task, which are very localised, which are relevant only to geographically bound communities, for very special and local interest in topics.

So for example, here in the UK, we have mySociety, a collective, a civic participation platform. With one example, such as FixMyStreet, being one of their projects, where they are enabling people to turn themselves into local reporters and report it to councils issues with their local neighbourhood.

Other examples are Smart Citizen, which is a platform enabling people who had an interest in the environment, to collectively gather data about, for example, air quality, sound pollution, and so on, and come together to share the data, build knowledge about their neighbour.



Cyclopath is yet another example of a community of cyclists, who want to promote of this mode of more sustainable travel. And they share their knowledge, in order to recommend each other cycle friendly routes, and also to suggest routes and this streets which require maintenance.

And, to me, these are all the incredibly successful examples of technology mediated collective intelligence, at least if we look in terms of numbers. So just going to pour some statistics, or some recent numbers that I'm aware of, if we look at the number of people that are contributing to Wikipedia, we have over 30 million user accounts. If we look at the number of people that they have been contributing to building OpenStreetMap, we have over 2 million user accounts.

So certainly, technology has scaled up participation, but then, we need to ask yourself the question is the size all that matters? And there have been studies, plenty, and as we know now, that looking behind what's happening-- we know that for example, in Wikipedia, over 80% of the knowledge has been produced by less than 20% of people. When we look at OpenStreetMap, it's even more dramatic. And what we have is a minority of 10% of people have been contributing to over 90% of the map, as it exists today.

So then, we needed to start thinking, OK, but then if these are active contributors we have, is everybody being represented? So for example for Wikipedia, do we have all genders, as we were discussing before, contributing to this knowledge? Do we have all ethnicities? Do we have all religions? Do we have all political views? And then, a followup question is if we don't, then what is the impact on the knowledge that we're gathering?

And indeed, there have been studies that show, for example, looking at gender is just one example, in Wikipedia only 10% of contributors are female. And in OpenStreetMap, only 5% of contributors are female. And there are big impacts. And studies have revealed, for example, in Wikipedia that there are articles covering topics like abortion, or biographies of a female assigned is not being covered as well as others.

So now the question is well, have we done anything wrong with technology? Because I would argue that doesn't mean-- so we are offering, as technologists, are very democratic. I mean it's not the case, at least in the Western world, that more men than women have access to internet. And it's not the case that the men are more computer literate than women. So what have we done wrong there?

And quite interestingly, studies have started to explore these issues. For example, for Wikipedia, what they found is that the female contributors base was larger. But over time, it grew a culture of conflicts and that editorial worlds that women not comfortable with, and they've thrown them away. So we're building technology, and then over time, it's been appropriated in ways which may not feed everyone.

Now, this is an example of what I've observed in these global scale collective intelligence platforms, but that we are observing similar forms of biases, also a more localised one. And just to give you an example, I'll mentioned TaskRabbit. This is a platform a for neighbours helping neighbours. And the idea here is that you might have some skills and time that you may use to help your neighbours running errands. So if you need a job being done, you post it



on TaskRabbit, and then someone else in your neighbourhood may pick it up and help you achieve it.

And over a relatively short period of time, just a matter of a few years, it's been picked up massively. So there are over 20 cities around the world deploying the platform, with over 25,000 workers, helping out over a million users.

But then again, the platform, per se, might be very democratic, but how is it going to be appropriated? And there was a study published in March of this year, where they have studied the adoption of TaskRabbit in Chicago. And what they have found is that workers are not uniformly distributed around the cities. And in particular, they are concentrated in the more affluent areas, with large underrepresentation in the economically deprived areas.

And this is an important impact, because what these researchers have found is that even when controlling by job type, and even when controlling by the distance that the worker has to travel, what happens is that workers are much less willing to go and do an errand in economically deprived areas. And if they're willing to do so, they charge more. As if to say, it's expensive to be poor.

So we need to be very mindful of the fact that, as technologists, we might put out a platform that is open to everyone, and people can take part in, but in practise, the adoption of this technology may have strong forms of biases. And with these forms of bias might change or overtime.

So what can we technologists-- we computer scientist can do as a way to help address the issues? And I'd like to offer a couple of examples from research we do in my group. So the first example, first thing I think we can do is we shouldn't only be putting out technology, but also we should put out a means of reflecting, and measuring, and quantifying over time how technology's being adopted, how we [INAUDIBLE], and what are the consequences on the task behind the technology we are putting out?

And I'll try to give you some examples from OpenStreetMap, which is a platform I've been studying for quite a few years now. So how many of you are familiar with OpenStreetMap? About half or maybe bit more.

So as I was mentioning before, you can think of OpenStreetMap as the Wikipedia of maps. Anyone within this room can go online and add a pub to the map that is being shown right here, which is an [INAUDIBLE] around Nesta, and it will immediately appeared there. However, what we know is that only one in 10 people here will have contributed information there.

So interesting questions for us are, for example, well imagine that in an hour time, we will close here, and we want to find the pub to go and have a drink together. So one question we may ask is are all the pubs which exist in the physical world being mapped there? The issue of occupancy we were discussing before. Maybe that doesn't even matter.

But let's say that we want to understand whether the 10% of people in this room, who have mapped most of the information being there, have put the same information that is of interest to the remaining 90% of us. So how aligned are these views?



Now the beauty, when technology's mediating this process is all about collective intelligence, is that every time we contribute a bit of information we leave a signal in the system, which remains there. And then, we can mine and study with data analysis process over time.

So in the case of OpenStreetMap, for example, we have a history from 2006 until '09 of every action that was ever done by any contributor on the map. So it enables us to run experiments and quantify the dynamics which are taking place behind the map.

Just as an example of work we have done, we have started with the case of London, for example, which is divided here. It needs 600 words or regions. And one question we started asking is how much knowledge exists in the virtual map that also exists in the physical world? And how complete is this knowledge?

So the first hypothesis that [INAUDIBLE] come to mind is, while areas close up to the centre would be denser and better covered. And also areas which have a higher population density will be better covered, simply because there are higher chances of there being active mappers there. But this doesn't explain at all. And, for example, what we have found in our studies is that areas which are, again, more economically deprived-- at least as captured by the Index of Multiple Deprivation in this country, the IMD-- are less covered.

And then we need to ask well what is the impact on the map? And we need, of course, to bear in mind that collective intelligence is not a state, but it's a process. So we may wonder, OK, if we are in that state today, where were we last year? Where are we going to be in a year time?

So what we have done is we have studied the evolution of the map over time. For example, here you see density maps from 2008 for a few years. And as you see, the growth is not from the centre to the periphery. And what we've tried to do, as a scientist, is build a computation on models that mimic these phenomenon that can be used in order to predict what is going to happen next. And though our models were fairly actual in leveraging history, in order to predict what areas would be covered next, and which not, by leveraging principles of spatial diffusion, of preferential attachments, and self-reinforcement.

So now, the point is are these tools useful? And I certainly think they are. Because, for example, what we can then do is say, OK, based on what we know so far, where do we predict we'll be next year? And then we can, for example, identify it as in London, the white spots there, which are areas that next year that will not grow. OK? Does it matter? Do we want to have those areas covered? Do we want to get knowledge about them? And if we want, then we can have technology interventions.

So for example, what if when I now go on to OpenStreetMap, the tool highlights what are the areas with that missing information? Will I go there and map?

For some localised communities, like Cyclopath, that worked. But these are small communities with high social capital. Would it work at global scale? We don't know.

But sometimes-- I'm not naive. I don't think that technology can solve all issues, but knowing that these are the areas with problems, and we may actually have offline interventions. So opens up, for examples, once weekly or monthly mapping parties, where people actually, in the physical world, get together, they start with a drink, and then they go and map the area.



These are the kind of offline interventions which might be informed by results of this kind of data analysis.

Now this was a case where we were looking, broadly speaking, at how much information was being mapped. But we also had more than once the questions. So we conducted a study where we compare data from 40 different countries around the world, and in this case, what we were interested in studying was to what extent the knowledge that was produced by this active community of 10% of contributors was aligned with what the remaining 90% of contributors were contributing. To see whether, for example, the same type of information was being presented or not, and using the same processes.

And these are some of the findings. Some surprising that we found. So what you see in this plot is every red dot is a topic that is being covered on open stream, from restaurants, to school, hospitals, and so. On the x-axis, you will see how much effort the crowd is devoting to mapping these points. And on the y-axis, you will see how much effort these 10% power users are investing in markings those items.

So if there is no bias, if the two communities, the two groups are equally interested on those topics, then the topic will sit on the line. And for us, it was a very surprising finding. So with the exception of one or two outliers, the two communities, on a per country basis, they mapped the very same thing. OK? So there's no need to change anything. The technology seems to be working fine, in that case. It's not being appropriated in different ways.

But what we found quite interesting was that if we then look at the way people are collaborating with each other, it changes a lot. So power users are behaving pretty much in the same way across all countries. But when you look at the other 90% of users, there are some countries, especially those that we found to have, for example, high power distance and low level of individualism, they don't override each other work.

So you can add an elements to the map, and then you might refine it. Or you may fix for spelling mistakes. You may fix geographical accuracy. You may add information. For example, to a restaurant you may add information about its type, its opening hours, and so on. In some countries, these are overriding or enrichment effects, they just weren't there.

So why is it interesting? Because, as technologies, the interface we provide is the same whether you're based in Italy, in the UK, in Brazil, in Japan, in the US, and yet, depending on the countries from which are entering your country are behaving in very different ways. So shouldn't we technologists be mindful of who is our user base? And shouldn't we adopt the technology to them?

So this is one line of study that I think, as technologists, you should be able to do. The data traces are there for us the mine. Let's build tools and [INAUDIBLE] that over time, continuously, will quantify adoption, and biases, and that we make these findings visible to the communities, so that we can either change the tool and experiment with new ways of broadening and opening up participation. But likewise, we think of offline and other ways of intervening.

And so-- do I've got another couple of minutes? Sorry.



No, I was just breathing.

No, I'm aware of time. So-- OK, so the other example I would like to give is on a completely different track. And my point of view is that-- I'm now moving more from global scale examples, like OpenStreetMap, to more localised example of engagement and participations.

And my view is that it doesn't matter where internet penetration we take us. For some of us, perhaps even for most of us, it will not be the case that we'll make the effort of opening up my laptop, going online into one of these platforms, and start contributing. It's too much of an effort. And the people and we will engage in some of these processes are already engaged once.

But cases where there is really something stronger, as [INAUDIBLE] was mentioning before, starting from the community, big issues that is throwing people together. So as technologists, we should think of live happening offline.

And for a-- sorry I'll just skip a couple of things-- and we should really think of ways in which we want people to engage and offer their opinions, voting, express, answering questions, which doesn't require them to make a conscious effort of going online and entering that information. But simply, that we stumble upon technology, as they go about their daily life.

And just to give you an example, there was Kassie in the room before, but she might have left. So at the beginning of January, we started a project where we were rethinking the UK census, so this data collection process that takes place ever 10 years, and that is very lengthy and tedious, and people are slightly disengaged with it.

So we started thinking, well is there another way of engaging citizens in answering questions of civic and even perhaps personal nature? Started discussing things. What these behind these data collection process and so on? And so we thought, OK, well why don't we beat the technology of spreading the Internet of Things, and we'll make it tangible and physical. And we will place it in the built environment, where people are already going.

There was no priming. It was never the case of grabbing people from the street and asking them to come here and answer these questions. And instead, what we did was we left the technology there, and people visiting the cultural centre may, on their own terms, in their own time, decide whether to go and interact with it or not.

And what they observed was that over 800 people interacted with the system, of the demographics, from young teenagers to elderly. Opening groups and often starting discussions around the question that were being asked there.

And then as another example, we tried something similar to get the opinions from children. It's very difficult to get the opinions from them, yet it's very important to understand what they think, whether the events and the things we are putting up for them are successful or not. And usually, you can only ask the question via the mediation of parents. Sometimes that's not what we want. So can we build tangible play for technology, that we put out in the environment, in the built environment, and that has some playful and physical element to it that we'll engage a broader audience?



So I think, that to conclude, I'm a technologist. I'm-- of course, technology is the answer the questions, but what is the question that we're asking ourselves? And I think that as technologists, we have to be particularly mindful that everything we put out there might induce a bias. But we have the power to make these biases visible and quantifiable, so that we can actually take informed actions. And I think that is our responsibility not only to put the technology out there, but also the means to reflect and make these biases feasible.

And web, social media, internet that was one thing. There's plenty more that is coming out way. We should experiment more with more participatory forms of action, broadening up participation. Thinking of where people already go, whether than expecting them to make the effort of going online and participating. And that's all.

[APPLAUSE]

My colleagues and I are trying to figure out what are the possible ways in which we can conceive of knowledge has been extended to the artefacts the we or, or even distributed between several individuals the same time. And today, I just want to talk to you about the second sense of knowledge being distributed, and how it might be of relevance to web science. And towards the end, I'm going to give you an example of how it can help us solve the problem that Wikipedia is currently facing.

So to start with, the general, philosophical question that I'm interested in is whether groups can acquire knowledge in a way that is functionality similar to the way that we acquired knowledge within our own heads. And so, I want to be clear, I'm not interested in knowledge as being the sum-- in group knowledge as being the sum total of the knowledge process by the individual members of the group, but instead as knowledge produced by the group as a whole.

So to give you an idea of what I mean by that, the starting point of my research is an interesting hypotheses within philosophy of mind and cognitive science, known as a distributed cognition hypotheses. Which says that sometimes when groups come together, under the appropriate conditions, they might give rise to an overall distributed cognitive system that consists of all of them. And it can do things that the individual members just couldn't do on their own.

And there are several ways that we can try and make sense of these interesting hypotheses. But I think the most interesting and promising way for making sense of it is in terms is dynamical system theory, which is the branch of mathematics that Newton introduced, in order to account for his mechanics, and which is used all over the natural sciences.

And the reason why I think that this kind of mathematics can be very helpful for understanding the idea of distributed cognition is because it says that if we have several distinct components, such as the individual members of the group, and we want to claim that these distinct components give us an overall integrated system that consists of all of them at the same time, like the distributed cognitive system, what we need to do is to make sure that all these distinct components are non-linearly related with each other on the basis of mutual interactions.



Because mutual interactions, first of all, give rise to certain new systemic properties that do not belong to any of the underlying sub-components, but instead on the ongoing interactions. So if you want to account for these properties, we will have to postulate the overall system. And secondly, also, these non-linear relations, they make it impossible to decompose the overall system, in terms of distinct inputs and outputs from the one sub-component to the output. The reason being that the behavioural of each sub-component is simultaneously dependent on the behaviour of the other components. And so if we want to make sense why those components behave the way they do, then, again, we won't be able but postulate the overall system.

So according to dynamical systems theory, if we want to claim that we have an integrated distributed system, what we need to do is to make sure that the individual members mutually interact with each other. And to see what this means in practise, it means that we cannot really go on and claim that every time we ask for directions from stranger in the street, we give rise to distributed cognition system.

Exactly because, in such cases, information flows only one way, and the cognitive processes of each individual are not mutually interdependent on the cognitive processes of the other individual. And so such a case won't qualify as a case of distributed cognition, exactly because the criteria of mutual interaction is just not satisfied.

And similarly, neither will cases of testimony, for instance, in the court of law, because, again, the information flows only one way. And there are no mutual interactions between the cognitive processes of the underlying individuals.

However, there are other cases of groups where the individual members mutually interact and would qualify as cases of distributed cognition. And just to give you a few examples, think about cases of brainstorming, scientific research teams performing experiments, jazz bands improvising, and I also think there are also a few other cases that, even though they don't include human components, they would still qualify as cases called of distributed cognition. And here, have in mind, cases of swarm intelligence, which is also where the idea of distributed cognition originates from.

Now, to see well this is related to the concept of knowledge, the connecting point is a longstanding epistemological problem, with a traditional account of knowledge as justified through belief, and specifically, a problem with a justification component. Because even though most epistemologists would like to say that justification is some form of ability to provide explicit positive reasons in favour of our beliefs, there seems to be lots of belief forming processes that generate knowledge, but no one knows how they work or why they're reliable.

And just to give you two obvious examples, think about visual perception or memory, both of which are supposed to generate knowledge, but no one, not even scientists, really know how they work or why they're reliable. And so we couldn't really provide an explicit positive reasons in their support.

So that's a long-standing problem that has puzzled philosophers. But recently, it has been suggested that the way to perhaps solve this problem is to give up this strong understanding of justification, in terms of being able to provide explicit positive reasons in favour of beliefs.





And instead, think about justification, in terms of cognitive integration, which just as in the case of philosophy of mind and cognitive science, so within epistemology, it is supposed to be a function of cooperation and interaction, or corporate interaction, with other aspects of the cognitive system. In other words, there is a striking similarity between the way philosophers of cognitive science and epistemologists think about the idea of cognitive integration.

But what is also very interesting about this approach, justifications, is that it explains the sense in which we can be justified on the basis of processes that we cannot offer an explicit reasons in their support. Because if a belief forming process, like our vision or our memory, is interconnected in this way with the rest of the cognitive system, it means that it can be continuously monitoring in the background. Says that if there is something wrong with it, then the agent will be able to spot these and respond appropriately. Otherwise, if there's nothing wrong, if the agent can be by default justified in employing the process and accepting its results, even if he lacks absolutely an explicit reasons to offer in its support.

So this is how thinking about justification, in terms of cognitive integration, solves a problem, that I mentioned before. But given also the striking similarity between the way philosophers of mind, and cognitive scientists, and epistemologists think about the idea of cognitive integration, it seems that we can then go on and claim that justification can, in fact, be distributed between several individuals at the same time, provided that those individuals give rise to an overall distributed cognitive system, on the basis of processes of mutual interactions.

In other words, if the presence of mutual interactions between the members of a group is what is required in order for a group to count as cognitive integrated, and thereby, also epistemically justified, then it does seem that we can go on and claim that groups can give rise to epistemic groups agents, that can generate group knowledge, in the sense of giving rise to true beliefs, that out of the product of a collective belief forming process, that arises out of socio-epistemic interactions.

Now, you probably think that this is quite vague, and such an idea doesn't have many real life applications, but actually this is not true. There are several ethnographers and philosophers of science who have suggested that the best way to understand knowledge produced on the basis of scientific research themes is in terms of distributed cognition, and something very similar to what I have just called group knowledge.

And another cool example also comes from cognitive psychology, where people are trying to figure out, or they discuss about transactive memory systems, which are groups of two or more individuals that collaboratively store, and code, and retrieve information. And the typical example is about an old couple. That we ask them, where did they get a souvenir from? And one partner says that it must have been more than 20 years ago. Which then makes the other partner say that it must have been during their honeymoon. Which then makes the other partners say something else, and so on, and so forth. Until one of them, or sometimes it happens both of them at the same time, make the finally recollection. And the idea is that on the basis of those transaction communication processes, the diet manages to navigate a said memory trace, and recall a piece of information that none of them could have recalled where they to act on their own.



So these are two complete examples of what I'd like to call group knowledge. And one more reason why I wanted to focus on those two specific examples is because in both cases, it has been suggested that what is practically necessary in order to have a well-integrated system is first of all, that every member of the group possesses some common knowledge that will allow them to communicate and start a relationship, even as strangers.

Which will then allow them to take the second step, which is to grow the differentiated structure of their community, by revealing information about themselves. And so allowing everyone to know who they are, such that everyone will know when it's time for them to rely on the knowledge and expertise of the other members, and when it's time for them to take action themselves. Which finally, will allow them to start interacting in mutual ways efficiently with each other. Which according to dynamical systems theory, is the only theoretically necessary and sufficient condition in order to have a web integrated system.

So these are two concrete examples that also demonstrate what is practically necessary to have a well-integrated system. But I think that the most interesting example of such an epistemic group agent is probably going to come from web science, where people talk about what they call social machines. Which according to Tim Berners-Lee, who's the father of the web, and the first to come up with an interesting idea, there's supposed to be web-driven processes, in which the people do the creative work and the machine does the administration. And which will enable us to just do things that we couldn't do before.

Specifically, says Berners-Lee, the idea is that if we manage to efficiently design such social machines, then high level activities, such as knowledge and justification, which have occurred just within one human's brain, will occur among even larger, more interconnected groups of people, acting as if they said a large intuitive brain. And so they do seem to be the kind of socio-technical systems that could give rise to what they would like to call group knowledge.

Now, you might think that this is some form of science-fiction, but so far, in fact, there have been several attempts to beat such social machines. And the most interesting and well-known case is no other than the case of Wikipedia. Which, according to a recent study, between 2004 and 2007, had an exponential growth in the number of its active contributors who are the editors that did not edit Wikipedia just once, but after the first time, they kept coming back. And they're also now thought to be the driving force behind Wikipedia.

Now, apart from being the most well-known case of a social machine, one more reason why I wanted to focus on this case is because it's currently facing a problem that I think can be resolved by thinking on the basis of distributed cognition and group knowledge. So the issue of the problem is, ever since 2007, Wikipedia has been facing a worrisome, steady decline in the number of its active contributors.

And as the rest of the study that I focus on indicates-- and they don't have the time to go through the details, but you can ask later on-- the reason for this is that around 2006, in order to keep the quality of its massive content high, Wikipedia introduced some fully-automated bots and semi-automated algorithms in order to prevent itself from being vandalised.

And the irony is that even though these fully-automated bots and semi-automated algorithms did help Wikipedia against vandalism-- they can actually spot vandalism within a few



seconds from the moment that it is been posted-- the problem is that they're also the reason why Wikipedia is now facing the steady decline in the number of its active contributors, because these semi-automated bots and fully-automated algorithms, they're very strict. And so they reject any entry that does not comply with all of Wikipedia's specifications.

But crucially, they do so without providing any feedback. And this is very problematic in the case of newcomers, because they see their entities being detected by default, without receiving any explanation as to what went wrong so that they can revise their entries to finally see them published. And so they feel as if they lost the time, and they never come back to become active contributors.

So that's a problem that Wikipedia is facing. And, obviously, in order to solve this problem, Wikipedia must find a way to keep the quality of its massive content high, without just using robots. And again, obviously, I think that the way for Wikipedia to do this is to find a way to allow its active contributors to meaningfully interact with each other by providing feedback, which is one of the general lessons that we take from dynamical systems theory if we want to have a well-integrated system.

And Wikipedia may try to do this in several ways. But I think that it can also use the help of the two main points that I mentioned before, in order to have-- which are practically necessary in order to have an integrated system-- which is to have common knowledge and differentiate structure implemented in its bio-technologically hybrid algorithm.

And again, Wikipedia could try and do that in a multitude of ways. But just to give you an example, and I will finish with that, one way that Wikipedia could do so is to first ask every existing and new contributor to also register some areas of expertise. Of course without asking for any credentials, because that would go against the completely free ability policy of Wikipedia, which has been so far very conducive to exponential growth.

And then, Wikipedia could use this information in order to more efficiently allocate the workload editing Wikipedia, by, for instance, sending notifications of a new entry to only those contributors who possess the relevant expertise.

But at the same time, Wikipedia can also keep monitoring how many changes a given contributor's edits undergo over time on a given domain. And if it is too many of them, recall that editor's status of expertise on the relevant problematic domain. And if Wikipedia manages to do that, then not only will it have managed to allow the right contributors to meaningfully interact with each other by providing feedback, but as the feedback in that workflow indicates, it will also have managed to epistemically self-regulate, by allowing the right contributors to keep monitoring each other's work and the stages of expertise, and thereby, what they can bring into the Wikipedia software. Such that, if there is falsehood or a mistake posted, then it will be swiftly spotted and almost immediately removed.

Otherwise, whatsoever stays online can count by default as justified and knowledge conducive. And if you remember what I was saying about how individual justification arises, this seems to be functionally similar to the way that we generate knowledge within our own heads. Thank you.



Just to kick off the discussion, a question for a Orestis and possibly for Licia, as well. Orestis, I'm very sympathetic with the idea of group knowledge, especially on theoretical grounds. And you know when you were talking, I was just thinking about the sort of question that Jeff posed before, which is a big challenge for all of us working on forms of collective intelligence, in other words, what sort of theory do we need, should we need, in order to explain collective intelligence?

It seems to me that you are going for a sort of top-down theory. So you focused mostly on distributed cognition, group knowledge, things that groups can do on their own, as if we ascribe agency and mentality to them. And then, of course, you go down to the individual level. And at that level, you have to postulate things like common knowledge, that are notoriously controversial, in many respects.

So I wonder if you have any thoughts on taking the opposite way and going bottom up? I mean, it seems to me that what you're saying about the distributed cognition is so distributed that the individual gets lost in it. And we are still trying to explain how forms of collective intelligence emerge from the interaction of individuals. So do you have any thought?

Yeah, I think that's probably not right, what you said about my approach. I think that I'm trying to do both. In a sense, my general approach is let's see how the individual mind works. And see whether we can functionally mimic that at the group level.

You know Minsky's idea of the Society of the Mind? That the mind is nothing else other than stupid micro agents. But if you make them intact in the right way, intelligence emerges at the individual level. So if that's how we should be thinking about individual intelligence, and it seems to get a lot of traction right now within philosophy of mind and cognitive science, then it seems that we can apply the same idea at the group level.

Now, I think that focusing on what the individual members of the group do is very important. So as I mentioned, they must have some kind of common knowledge. But that is not a very controversial ideas. It's mostly-- for instance, if we are just brainstorming, and we are a causal team, we need to speak the same language, you know? If we are scientists, we must have the same paradigm, what Thomas Kuhn was talking about.

But then at the same time we must not just have that common knowledge. We must also have a differentiate structure. Every person must have a specialisation. So there are things that are really important about the individuals for generating group level completion.

To Licia, is that right? Yeah. I was very interested in your argument about reflecting on online platforms and so on. One of the things that worries me most is the digital divide that is now being created between who actually has the tools, and the power, and the money to actually analyse this data, from universities to political parties, and so on.

So it was questions that I wanted to bring up also this morning. Like what happens if these kind of mappings end up in the wrong hands?

The first thing you seen is that now we have tools that anyone in this room can use in order to build the next collective intelligence platforms. So maybe 10 years ago it took lots of skills in



order to build the next Wikipedia. Now anyone within the room could build an example of crowdsourcing platform or collective intelligence platform in a matter of a couple of hours.

Now, the next wave and a bit behind that is a lot of academics have said, well actually we needed to quantify the processes which are happening behind, and make them visible. And quantified forms of biases either in the contributors base or in the outcome of the task. And what they're observing, actually, is that we are also putting out these as tools that anyone can then plug them in into their websites.

So this is exactly what I'm advocating for. So we need to make democratic access to the tools to analyse and reflect upon what's going behind the scenes. So we're not there, but this is exactly what we're pushing for.

I have a question, which I would like you to answer, Orestis, and his plea for-- is it something conceivable what he's arguing for in Wikipedia terms?

Eh, sorry start again. I was lost in other thoughts.

Orestis finished off with a sort of plea, a way of fixing Wikipedia. And so I'm asking you, is it justifiable to use--

So first, I'd like to add an example, and then I'll go back to the question. So what you have observed in Wikipedia, that idea that all of a sudden you have these automation intervening and changing the dynamics of groups, is actually not constrained to Wikipedia. And there have been plenty of other examples, which I think what might be interesting in.

For example, in OpenStreetMap, example that you have seen is that if you have a blank map, and you have to start editing, that might be quite an interesting for some people, because if you know I've got power to have knowledge there. But for other, it might be, oh, this is an empty map. It will take me ages to go there.

So in some countries, what that has happened where actually mapping data was there, why don't we just import it if the agencies who own it can give it to us. So we have these bots, which in matter of an overnight were filling up the mouse with lots of knowledge.

And what were the effects? In some communities, in some countries, it spiked contributions. Because they said, oh, there's already a note up there, so I can actually start using it quickly. But in other countries, it actually completely destroyed efforts, because it said well, if you can do it for free, why should I make an effort?

So to me, the point is, it's not going to be the case that it can be all automated or all human. I would even argue that it's not the case that we need humans in order to check on the works of others. It's a matter of understanding how we can actually collaborate between humans and robots or bots. So the dynamics should then be human to human, or accepting the work of robots, but more of an interactive nature between the two.

And we still don't know how to do that. And even in the case-- I'm particularly interested in geographic knowledge. And even this case, we are observing new scenarios of operations, especially if the area of human Italian mapping, in response to disasters.



Where on one hand we have these drones are flying around capturing imaging information about the environment, but this is very stellar information, so you have mapping information with no names of what the villages were there. And you need domain knowledge to decorate on top of that.

And this is a case where you need interactions between machine contributed data and domain knowledge. And again, we don't know what the relationship and dynamics between the two are. So I think what we need to study next-- and even as technologist, we need to enable next are exchanges between automations and human efforts.

Just a second. When I-- to make a link between Paulo and Licia-- because there's something that you didn't discuss that is in your paper.

By the way, I would like to make the apprentices that we have a drop box. I cannot give the drop full of intellectual property problems to everybody. But if people want to write to me, I may give them access to the drop box file, which has a lot of the papers that are being discussed.

So in your paper, you talk about the effect, in a certain sense, of technology on the choice of form of deliberation. And I'm just wondering if there's something to be discussed? Maybe you want to mention what you describe [INAUDIBLE] in the paper, and the effect that these different forms are having on forms of deliberation with respect to what Licia just said?

Mm-hmm. Yeah.

I'm sorry.

Licia was pointing out that in a way there is a sort of gap between technology and then the social effects, right? In a sense that you can have the most democratic technology ever, and then, somehow it would always reproduce some biases in society, right?

And I think it's true. There is somehow an element of neutrality of technology, in that it could be used for very different purposes. And somehow, there are some limits to it, in terms of its social impact.

At the same time, it is also seen that technology often can be kind of politically, very [INAUDIBLE] political informed, in that some of the decisions that apparently are kind of technical decisions, as to the interface, as to, for example, the complexity of the interface, for example, are ultimately kind of political decisions, masked as technical decisions, right? So that's why it's often very dangerous to get designers to make decisions, because they might make something mistake design decisions for what are actually kind of political decisions.

I mean, just imagine a platform like OpenStreetMap that can be made more user friendly, but then would mean less interactive, possibly, right? And so, in that case, it can appeal to people who are very technically skilled, or it can appeal basically to dummies, so to speak, right? And you see that in very different situations, the kind of pitching sort of the interface to a certain kind of category of people.



I think that's-- when I was talking to some of the people behind OpenStreetMap, the first time five or six years ago, I tried to do an edit. I've got a background in computer science. I would call myself as tech savvy. And then the first time I tried to make an edit in OpenStreetMap, I couldn't do it. So the technology to me was unbearably difficult.

And then when I was talking to the people behind it, they say, cartographers, they said, we did it on purpose, because we want to make sure that who edits the map has the cartographic skills to do that. So we embed quality. Somehow it was a way of filtering. Yeah

But now, OpenStreetMap is being used also as, not only for a cartographic knowledge, but also, for example, in humanitarian responses. And in projects like missing maps, where what you need is a disaster is up, and then in the course of 24 hours, you need to map whole countries, and there's nothing there. You need to know what the entry point, so that many more people can contribute to that.

But then, again, there's an interesting point. At the moment OpenStreetMap doesn't have any processes by way of controlling the information which is being uploaded on the map or not.

And what I've observed, when people use these interfaces in order to upload information as part of missing maps, for example, is that often they do not commit the edit. And we ask them why. And they said, I don't feel I have the authority to do that. I would like someone to actually check that the edits I've done are correct, because then there would be somebody from the British Red Cross out there in the field that will use my information to go and check that village.

So again, this is a case where perhaps having these tools embedded would help. So any choice we make, I totally agree in terms of--

Can I? Can I? Can I? Yeah, OK. So you applied the-- we could do away with a human components. But I just want to say, because it's all related to what you've been saying, that the idea of a social machine, which says that the human components must be doing the creative work and the machines should be doing the administration, shouldn't be taken extremely lightly. Because there are some tasks that, at least for the time being, machines couldn't conceivably do.

So for instance, deciding whether an entry is good one or not in Wikipedia is not something that the machine can really do. Maybe highlighting [INAUDIBLE] areas in a map have not been covered, that is something that a machine can do. However, in the case of the maps again, a machine couldn't give you any stats on whether a bot is a good one or not. You need a human component for that. So I don't think that we should we just try to do everything just using machines.

Oh, I totally agree.

Yeah, yeah, yeah, yeah.

I'm sorry. There was a question over there. Sorry, Lydia. You're running, but he--



Yes, my issue is related to the last two questions, in terms of the technology induce bias. Quote, people are creative whereby are machines are administrative. In that respect, should we really trust the outcomes we are given for collaboration, despite not knowing what goes on behind the scenes?

I do not understand the question.

So the question was should-- I'm sorry. They didn't really-- could you say it again?

Yeah. I have an issue with the issue called the technology induced bias. To quote, people are creative. Machines are administrative. In that respect, should we really trust the outcomes we given, in terms of research and collaboration, not knowing what goes on behind the scenes, in terms of what happens to research, the information, et cetera?

I think your answer is about providing that sort of transparency.

I'm not really sure what you mean by humans are-- OK, I understand what you're saying about humans are creating and machines are not, but in relation to technology bias, I'm not really sure what is the point you want to highlight? What I was trying to say in terms of technology induced bias is that even if the technology we put out there in principle or access to it, even if in principle, it might be democratic, meaning that anyone around the room can access it. That we make a decision. We self-select whether to take part or not, and the dynamics with which we participate, even just in terms of how much, we collaborate online. What language we use when we did we discuss opinions in a forum an so on.

They all are driven by-- they vary and they are appropriated in different ways by different people. And these might induce some people, for example, to leave the platform. Some people might say these communities are way too conflict prone. I don't feel comfortable with that. I might leave.

So when I'm saying technology induced bias is that there are-- people appropriate the technology in very different ways, which might then have an effect on who decides to take part and how.

[INAUDIBLE] the outcome. My point is that there's not enough visibility of what is happening. There is not happy visibility of who talking part, how much they're taking part, what they're doing. And in order to build trust in the outcome, I think we should build more transparency in the processes that are leading to these outcomes.

