# AI and You

Transcript

## Guest: Louis Rosenberg, part 1

## Episode 158

First Aired: Monday, June 26, 2023

Hello, and welcome to episode 158! What do honeybees have to teach us about AI? You'll find out from my guest today who is Louis Rosenberg. He has done some just fascinating things, starting over 30 years ago in the virtual reality labs at Stanford and NASA. In 1992 he developed the first mixed reality system at Air Force Research Laboratory. In 2004 he founded the early AR company Outland Research which was acquired by Google in 2011. He received a PhD from Stanford, was a tenured professor at California State University, and has been awarded over 300 patents. He's currently CEO and Chief Scientist of Unanimous AI, a company that amplifies group intelligence using AI technology based on the biological principle of Swarm Intelligence, which is where the bees come in. The UNU Swarm Intelligence that he created has an extraordinary record of predicting, say, the winners of the Oscars. Let's get right to the interview.

Louis, welcome to AI and You.

Yeah, thanks for having me.

You've done fascinating work and I want to talk about the metaverse and AI for its potential and its peril with respect to shaping our beliefs and manipulation. But I want to make sure that I talk especially about hybrid swarm intelligence and your fascinating TEDx talk about honeybees and what you learned from them and what you've built with them in UNU. So rather than me going on about that, why don't you introduce us, for people who don't know what UNU and hybrid swarm intelligences are, to that.

Yeah. So really it was about a decade ago when I started thinking about, "How can we use AI to enable groups of people to make better decisions, smarter decisions, wiser decisions; decisions that are better aligned with the whole population?" And like a lot of technologists, I look to Mother Nature for inspiration and guidance. And it turns out that nature has been wrestling with this problem for hundreds of millions of years, and biologists call it swarm intelligence. And so, swarm intelligence is a property that's found across many different organisms where you have large groups and those groups don't necessarily have anybody in charge but need to make decisions that are good for the whole population. And so, swarm intelligence, the reason why birds flock and fish school and bees swarm, they become significantly smarter together than alone. And so, if you think of a school of fish, let's say 1,000 different individuals, each of those individuals has a slightly different view of the world. Each of those individuals has a slightly different temperament, slightly different history. Nobody's in charge, there's no leader, there's no hierarchy. And yet somehow together they can function as a superorganism and they can navigate the ocean and make all the decisions they need to do as a superorganism to survive. And they've survived for hundreds of millions of years. Swarms of bees also make decisions in large

groups, and birds. And so the question that hit me, and this was again about a decade ago, is, well, if birds and bees and fish can get smarter together, could we use the same techniques to allow human groups to get smarter?

So if the birds and the bees can do it, why can't we?

Right. And especially since we've built this infrastructure for connecting us all together in real time over the internet, can we enable similar techniques? It turns out that birds and bees and fish have evolved very special techniques for doing this, but what they do that is really very different than humans - so, us humans, we exist in large groups, we need to make decisions in large groups – and what we tend to do is take a vote if it's formal, or maybe we'll just take a poll of a large group, and from a poll find out what are the sentiments in the group and maybe use that to inform decisions. It turns out that nature does something really different. It doesn't take votes or polls or surveys. What nature does is, it forms real time systems, real time systems where the groups can actually push and pull on each other and interactively, dynamically, find the solutions that they can best agree upon. And so a school of fish does this because each fish in a school, they can detect subtle vibrations in the water around them. And those vibrations inform them about the sentiments of the fish around them. And these are all overlapping sentiments. And so you end up with this superorganism. And so what we did - and going back to 2014 - we said, "Let's build a system." And the first system we built was called UNU. We then built our current systems called Swarm, but they work similarly, but our technology called Swarm enables groups from anywhere in the world to log in and connect in real time. And because we humans didn't evolve to detect vibrations in the water around us, we use AI as the mediating technology to allow the groups to combine their sentiments. And so in the swarm system, a question appears on everybody's screen at once. Each person has an interface. They can just use their mouse, their touch screen, and they can use what looks like a little graphical magnet to pull on the group, to pull on the swarm. And so you could have a hundred people or a thousand people. A question appears. Everybody's pulling in different directions. And it becomes this multidirectional tug of war while real time AI algorithms are watching everybody's behaviors. And so our system is trained to look at how do people behave, to infer their different levels of sentiment. And so it infers who is confident in certain answers, who's ambivalent, who conceded to the group very quickly, who is entrenched in their views. And it's a real time system. So, what happens is a question appears, this whole swarm of people behaves instantly to the question. The AI watches how everybody behaves and guides the swarm in a direction towards an answer. But as soon as the swarm starts moving to an answer, everybody changes their behavior. And then the AI looks at how did they change and has a better sense of their sentiments. And so, it becomes so the AI is influencing the people, the people are influencing the AI, and it basically helps this group converge on the answer that they can best agree upon. And it turns out that it works. Mother Nature was right. And it works for human groups. And so we've done over the last 10 years, we've done studies with MIT and Stanford and Oxford and Cambridge and Carnegie Mellon and where we've tested this idea of creating what really a real time hive mind of sorts where everybody is participating in real time. And we see even stronger amplifications of intelligence than we expected. So, for example, we did a study with researchers at Oxford looking at just groups of sports fans predicting soccer, right? In English, English football. And so there's just

groups of regular sports fans. And then they predicted 50 consecutive games in English Premier League. And when they predicted, you know, as individuals, they were OK. They're about 50 percent accurate, which is which isn't bad because it's win, lose or tie. When you did what people call Wisdom of Crowds, which is just taking a statistical average, they got a little bit better, 55 percent accurate. But when they work together as a system mediated by AI, they jumped all the way up to 70 percent accurate in predicting the outcome. And we've seen similar results with Stanford Medical School. We had groups of doctors, small groups, just five doctors making diagnoses. And we saw them reduce their errors by 26 percent. We had we did a study with MIT, with groups of financial traders predicting the price of gold and oil and the S&P 500. And they increased their accuracy by about 30 percent when they worked together as a swarm. And so it turns out that when you connect together human groups and you use AI to infer their sentiment, so we really can combine the strength of their convictions in a in an accurate way, these real time systems become smart and enable human groups to do to do something that they couldn't do on their own. And yeah, and so my philosophy is really about finding ways of using AI to amplify human intelligence while also keeping in the loop inherently human values and morals and sentiments. And so we're you know, we're allowing decision making. In ways that are not removing humans and our unique abilities, instead, we're using AI to connect humans.

There's a beautiful philosophy at the core of this. The attention grabber for me was that it predicted the Kentucky Derby Superfecta. I think the reporter made $500 following that one.

Yeah. So we one of the things that we've one of the things that we've done over the years is take on challenges from journalists. And so I'm saying, "Hey, you know, we can take a group of people, we can make them smarter". And journalists want us to prove it. And that was a case where it was a journalist from CBS Interactive. She said, "well, can you predict the Kentucky Derby?" And again, what's interesting is that we don't know anything about the Kentucky Derby. We don't we don't have an AI that's trained on horse course racing. What we have is an AI that's trained on human behaviors, how humans behave in a swarm. And so, we can take a group of doctors or we could take a group of racing enthusiasts. And it doesn't matter. They come into the system. And so, we predicted. So, the reporter had asked us to predict not just the winner of the Kentucky Derby, said, but give it, you know, predict the first four horses in order because it's called the Superfecta. And it was 540-1 odds. And so, it was like we said, "OK, well, we'll give it a shot". We brought in I think it was twenty-four racing enthusiasts. They came in, they made the prediction. We gave it to the reporter and she went to the Kentucky Derby and she wrote a story, which put pressure on us. But then she went to the Kentucky Derby. She placed a bet and she tweeted it out during the race. And anybody who placed a twenty-dollar bet would have won eleven thousand dollars. And I actually placed a twenty-dollar bet, so I won eleven thousand dollars. So that definitely generate a lot of attention for us. But these days we have more interesting examples like the United Nations has used our swarm technology for predicting famines around the world and they have groups of experts with very diverse backgrounds, experts who look at the economic conditions in different regions, look at the climate regions, experts that look at the economics, look at political. And so you have this issue of, you've got a group of experts, they're looking at what's the probability of that there's going to be a famine in Haiti over the next 18 months, how do you combine all their different sentiments? They you

know, they have different they have different expertise. Well, when they come together in a swarm, you basically become a super organism and they can converge on forecasts that are better than they would if they just took a vote or if they just argued about it.

There is a group of people, a type of person called Super Forecaster, and I wonder if you had the opportunity or would you like to get Super Forecasters as an UNU group, a swarm group and see what happens when you put them all together?

So, yeah, there are people called Super Forecasters. They've either inherently are good at forecasting or they've trained themselves to be very objective and trained themselves to be good at probabilistic forecasting. It's a it's a great technique. People can become better forecasters, it is a skill. And we've looked at when we build swarms of people, if we build a group, if we take a swarm of regular people, average consumers, novices, we can turn them into an expert. Or if we take a group of experts, we can turn them into a super expert. And so, the smarter or the better forecasting capabilities that a group has, the better it works to swarm. Because one of the things people often ask is, "What's a good group?" Really, the best the best way to build a hive mind or this type of artificial expert is to have a group of people who know what they know and know what they don't know and are well calibrated in their levels of conviction. And a Super Forecaster is really somebody who's good at doing that. There are other professionals like we we've worked with hedge funds where we have groups of financial traders. They tend to know when to be confident, when not to be confident. There are some groups of experts who are actually very bad at that. Like we did a study with Marist University because they do a lot of sports forecasting. And they said, "Well, let's have sports fans. Let's make a swarm of sports fans and let's have them compete against a swarm of sports writers". And so sports writers are considered to be these great experts and we think the experts would be better than the sports fans. It turns out the sports writers were very, very bad forecasters. And it's because they've trained themselves to write about the interesting things, the interesting angles; not the most likely angles, right? Like if they always wrote about what was most likely to happen, maybe people wouldn't read their stories.

Right.

And so it's sometimes experts are overconfident or they think they know more than they do. And sometimes not novices are really very well calibrated in terms of knowing. But as you said, super forecasters are people who really do study this issue and know and know when to be confident and when not to be confident.

I want to pick up the role of the AI here because the user interface for these sessions looks like an array of possible answers as vertices around a circle, of a polygon. And each user controls this little magnet, as you said, drawing a ball, which represents the group think, I guess, towards the answer that they want. And I would always assume that the ball moved according to like the vector sum of the forces on it from the magnets, and this was then proving out that principle of averaging of human evaluations, like the average number of guesses of the number of jelly beans in a jar at a state fair is closer to the answer than any individual. But it sounds from what you're saying as though it's more complicated than that. How is the ball moving?

Yeah. So the key thing about a swarm is that it's happening in real time. It's not it is. So, you know, a traditional say wisdom of crowd is I could I could ask if I have 100 people. I could ask them all to give me their prediction. What's you know, what's going to be the price of gold in two months from now, and they'll all give me a prediction and traditional Wisdom of Crowds says that, you know, the aggregation of, you know, let's say a thousand people will be significantly more accurate than any than any of the individuals or the vast majority of individuals. And so that works. The problem is that that aggregation is really has very little information in it. I know what everybody says, but I have no idea which people are really confident in their answer, which people are just guessing. I have no idea of their conviction. And to assume that all their convictions are the same is just be wrong because it's not how that population would be. Now, I could ask on a survey, "Tell me on a scale of one to 10 your level of conviction". Turns out people are terrible at doing that. People do not know their own level of conviction. And if they do tell me on a scale, what you say is a seven and what I say is a seven might mean completely different things. Our scales are different. Our scales are nonlinear. And so you have this fundamental problem. You have this this concept that crowds are smart, but we really don't have sufficient information about their conviction. When we bring the group together into a swarm, now they're all interacting in real time. And we're not just asking them to just click on an answer. We're asking them to behave. We're saying, "OK, you're going to behave in real time. And we're going to watch what you're doing to infer your level of conviction". And so the interface, you know, there's this like you said, there's this little round glass puck and each person has a magnet and they have to continuously behave, meaning they put their magnet next to the puck. And if the puck starts moving in their direction and again, there might be 100 other magnets, people pulling in other directions, but if the puck starts to move in their direction, they have to back up to get out of its way. And if the magnet starts moving away from them, they have to chase the puck and keep their magnet close or they will lose influence. And so they have to be continuously adjusting in real time to convey their impact. They also can see what everybody else is doing and that's influencing them as well. And so, what happens is, if the group is if everybody's pulling in different directions and let's say there's a stalemate, people won't just sit there in a stalemate. Somebody might say, "Yeah, I was pulling for this. You know, I have these different answers, but I was pulling for this one and maybe it's worth I'm totally fine pulling in this other direction, which maybe I had very similar opinions about". Somebody else might do the same thing. Somebody else might just see that the puck is moving in a direction and just think, "Well, the group probably knows more about this than I do". And they just concede. And so there's a lot of different things people can do. They can entrench, they can concede, they can switch to another option and they could do that right away, or they could they could resist for a while and then realize that it's hopeless to go to one option and they switch. And so the AI is watching each person's behavior and inferring their levels of conviction. And it's doing that every 200 milliseconds, and so it's updating everybody's conviction and then the actual swarm moves a little bit based on everybody's conviction. But as soon as it moves a little bit, that changes the state and everybody behaves a little differently and changes their conviction. And so what you end up with is with this this system where the puck might start moving in one direction and stop and then it'll basically find the path to the answer that usually optimizes what they can best agree upon with consideration of their different levels

of conviction. And when we interview people, we find they'll tell us they didn't even really know their level of conviction until they engaged in this process, right? They you know, if when you bring together, let's say, we did an NSF funded study at Stanford Medical School, we had radiologists. I mean, radiologists, they're experts. They've gone to school for 12 years to become a radiologist. And, you know, a chest X-ray pops up on five different screens of these radiologists, and now they're going to predict the probability that this patient has pneumonia. Well, you would think that they're all pretty confident in their view, but as soon as they start pulling for direction and they see that the puck isn't necessarily moving in the direction that they thought, it forces them to introspect on, you know, "Am I missing something? Did I, did I not?" And so and so they're just so the group that the system is both assessing their conviction.

Right.

But also evoking their conviction as they discover. And we study after study. We see that groups make really good decisions, forecasts. And I've talked about this in the context of, you know. Decisions like, you know, what's the probability of somebody has pneumonia looking at a chest X-ray, but it also works when you're trying to combine sentiments, meaning, if I have a business team and I'm trying to find which of these different projects do we think is the most important for us to work on next? There might not be a right answer, but there is an answer that would maximize their collective satisfaction, which is a little different than their collective conviction. Like, if it's a forecast, it's a collective conviction. If it's something that is softer, it could be their collective satisfaction. And what we find and we do studies with big brands looking at predicting marketing messages and predicting which products features are most important when a swarm converges on an answer, it is really finding the the solution that that really tells us what is the collective sentiment or the collective of this population, even for an answer that is more opinion based than forecast based.

That's fascinating. The AI then is really acting like a many to many real time sentiment multiplexer. It's like it's hooked up this set of channels between all of the members of the group that relays the how they feel about this issue to each other. That's amazing. Is there an analogy, a parallel between the members of a swarm and neurons in a brain, or is that misleading?

I think there is an analogy. I think, you know, like every brain analogy, it's loose, right? But I like to think of a swarm really as a brain of brains, and I say that because when we bring a group of people together into a swarm, the thing that we're doing that's really different than, say, taking a survey, when you take a survey. You're taking a person who has, you know, a lot of knowledge, wisdom, insight, and you're reducing all of that down to one or two data points. When you when you connect a group together in a swarm, we're treating them more as a data processor, right? We're saying, "OK, you. You're going to interact in real time, we're going to draw upon your memories and your intuition and your wisdom, and it's going to and we're going to help you evolve, help you help the decision evolve". And so we really do like to think of it as, instead of collecting data from people and statistically aggregating it, we're actually just connecting the people, and we're allowing the aggregation to emerge based on their interactions. Which is, again, very similar to, say, a brain, you know, brain is about, you know, you have this network of neurons, each one has different activation levels. They're all influencing each other.

And an answer emerges from these interactions. When we connect, we connect a group into a swarm, we're essentially taking it up one level of abstraction.

Right. And clearly it breaks down when you start thinking about how message passing works in among neurons. But I want to think about a bigger application of this, like if Wisdom of Crowds was accurate, like the averaging of many people's input was the best answer, then the outcome of the United States presidential election would be the person best qualified to serve the interests of the country for the next four years. I suggest that that hasn't happened in many cases. What would an election look like if it was done with this real time feedback?

So we do we do a lot of work in political polling, political forecasting using swarms, we tend to get very accurate answers compared to traditional polling. It's a really interesting topic because we often think of taking a vote or a poll as giving us a reflection of the population. It turns out that polling is a double-edged sword because it also influences the population. And so what you have, I often like to say, the fundamental problems of polls is that they're polarizing and they're polarizing because what it does is it tells us it shows us the differences in a group. It doesn't tell us the commonality in a group, doesn't tell us how that group is better able to find common ground. It really very sharply shows us the differences. And then unfortunately, we publish the polls and we actually force those differences to entrench where we. Our obsession with polling actually drives populations further and further apart and gets and drives more and more entrenchment. A swarm, because it's a real time process, is actually doing the exact opposite. If everybody's pulling in a different direction, it goes nowhere. It will only find a solution if the group can find where they best agree. And so it actually instead of amplifying differences, it amplifies commonality and it allows groups to discover, to discover this the solution that maybe nobody was nobody's first choice, but it's the solution that they could best agree upon. And it makes sense from a from an evolutionary perspective. You know, a school of fish will never do what humans do. A school of fish will never, you know, have multiple factions just entrench where they cannot agree what direction to go. You know, there's a there's a predator bearing down. The school of fish does not split, you know, entrench or split into two schools of fish that go off in different directions forever. Evolutionarily, probably all the biological algorithms that that that achieved such a school of fish died out and the schools of fish that that discovered mechanisms that allowed them to find the solution that they could best agree upon, survived. The same thing with swarms of bees. Swarms of bees have remarkably sophisticated decision-making processes, even more sophisticated than fish. One of the things they have to do every year is find a new home to move into, and it's this really complicated decision where they have to go out and they send hundreds of scout bees into the world to find potential home sites. And they actually bring that information back to the colony and then they might have a dozen different sites and they need to pick the site that is the best one to move into. And it turns out it's like this complex multivariable problem because they evaluate their sites based on, you know, is it does it large enough to store the honey they need for the winter? Is it insulated well enough to stay warm on cold nights? Is it ventilated well enough to cool to cool down in the summer? Is it near good sources of water? Is it near good sources of pollen? It's this really sophisticated decision. And biologists have studied these decisions over the last 50, 60 years. And it turns out that the bees almost always pick the optimal solution across this complex multivariable problem. And they do

it by forming this real time swarm where they're pushing and they're pulling in different directions. They actually do it by vibrating their bodies to generate like the different scout bees will vibrate their bodies to generate signals in support of different options. And when there's a stalemate, they'll switch their choices and they'll ultimately converge on the solution that they can best agree upon. And it is it is almost always solving this complex multivariable problem. And bees have their brain is smaller than a grain of sand. And yet it would be hard for a human to solve that problem. And it really is great evidence that that when you build this brain of brains, essentially, it becomes something much more than the individuals would be on their own. And what nature shows us is it's allowing the group to become this real time dynamic system.

It's a beautiful process. And I didn't appreciate its beauty from this standpoint, until you articulated it like that. And now it's that's so that's a wonderful gift that you've given me and others there.

That's the end of the first half of the interview; this one's split into two parts to be more manageable for your time. You can read about Super Forecasters, by the way, from the 2016 book *Superforecasting: The Art and Science of Prediction,* by Philip Tetlock and Dan Gardner. You can see UNU on the web at unanimous.ai, check out the videos under the resources menu.

By the way, this marks the completion of our third year of podcasting, and it only gets more fun and useful for me and many other people as we continue to hit record numbers of listeners. Please take a moment to give us a five-star review and a nice comment on your podcast platform, it makes a huge difference to our ability to reach new audiences of people – pretty much the only way that we reach new audiences - who want reliable, high-quality information and thought-provoking questions about AI, and more people like you.

In today's news ripped from the headlines about AI, Google is prototyping a system that lets robots write their own code to respond to instructions and perform tasks. They have open sourced the robot control method, called Code as Policies, which uses a large language model to generate code for the robot from commands like "Put these objects in a horizontal line near the top." The language they can understand is somewhat reminiscent of the very early system called SHRDLU but obviously far more capable. The fact that there is code generation and robots involved doesn't mean that there's potential for robots to run amok, in case you were wondering, that was just something that the headline writers liked.

Next week, we'll conclude the interview with Louis Rosenberg, when Louis will tell us more about the philosophy behind his work, the ideal roles of humans in cooperation with AI, some of the threats to privacy posed by AI, and the convergence of AI and AR/VR technology like Apple's new headset. That's next week on *AI and You.*

Until then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

**http://aiandyou.net**