AAAI/CCC Symposium on AI for Social Good

Talk Sessions 3: AI for Urban Planning
Session Chair: Dr. Virginia Dignum

Fei Fang: Today the first discussion will be on urban planning, and we are very fortunate to have Virginia Dignum to chair the session for us, and Virginia is associate professor at the Faculty of Technology, Policy and Management at the Delft University of Technology. She received her PhD in 2004 from the Utrecht University, and her research area is agent based models of organizations, and she's one of the leading researchers in this area, and she was the co-organizer of AAMAS 2005, and also the co-chair of European AI Conference in 2016.

And with that let's welcome Virginia for the opening talk. And by the way, I also want to thank Steven Smith, from CMU, for his support for this session, although he cannot make it to the symposium.

Virginia Dignum: Thank you. First I have to apologize as to make a note on the Delft blue of the slides, I think Delft will be very unhappy if this is being filmed, because we are very proud of our Delft blue, and it become green in this projector, so don't take the green literally, and we have very strong connection with Delft blue. Anyway. So this session is on urban planning, and urban planning is an increasingly important area for AI, and for the technology development. Already today, more than half of the population of the world is living in urban areas, and expectation is that two thirds of the population will be living in urban areas by 2045.

Oh, and that one is red, something is very wrong. Anyway. So, and this whole migration of people from rural areas to urban areas is meaning increasingly higher stress being put on the contemporary urban systems, and a lot needs to be done in order to keep cities being sustainable, and livable, in the coming years. So this is managing the urban areas, is one of the very important challenge for the 21st century, and one in which AI can do a lot.

AI and urban computing focus on application of intelligent computing technology to the problems faced in urban communities also tries develop solutions and technology that can support solving, or at least minimizing this type of problems in different key areas, and these areas are some of the areas in which the papers of the session of today are going to be focusing on, but I think before we really go into the papers, what I would like to convey in this short introduction is to my idea we need to, not only look at how to solve the problems, but also take a kind of a step back and think about why, and how, and what are we going to solve, and why should we solving that.

So, in a sense, what does it mean to have good AI or AI for the social good, applied to urban planning. One of the important issues is to look at who is planning. In many case ... that's very wrong with this slide, something didn't ...

Anyway these pink guys are supposed to be a bunch white guys sitting in some office and planning, and that I think is what is notion, that one girl I think in between, but they're all white. And this is a lot how urban planning is done nowadays. People with expertise in planning, they go and decide and make the decisions on how cities, and how urban environments look like.

There is, of course, a lot of going into participatory planning, which is the other one on each side should be saying something like, "What would you like to be built here in this area of city." It is very bad slide. And in this participatory, it's a step in a good direction, but a lot of it, what we notice is that people who are participating are representing very specific areas of the population, so they are typically the people who join these kind of discussions, are people with the high education, and more verbal, and more committed to the society where they live, and not all the lower and poorer areas. Poor people, or some of the neighborhoods on which they are a lot deprived people and populations living on.

One the issues with urban planning is that, literally, the decisions we are taking on the design of cities, are going to be set in stone, and by being set in stone, or in infrastructure, they will impact the life of the societies for a long, long time. And they are really making a big difference, the way we implement, or we decide in decisions for urban systems, it's really to stay there for a long time, so that's something which is very important to consider, before we really are going to determine what to do. And of course, we all have very good intentions, and I don't doubt that the pink guys there are all very well intentioned, and they really know what they are doing, but not always those good intentions lead this necessarily to a good design.

And there is much more to think about when we are taking those decisions. So examples of politics, which are ... or decisions, which really make an impact on the values of people, is those low hanging branches, which are often ... or there are all kinds of studies which show they are used kind of to prevent populations, from poor parts, to best is to go to the other parts of the city. You have those park benches in those pictures there, which you hardly can see, which are designed explicitly to avoid that people sleep on them, so it's some kind of avoiding homeless people to take a space in those parks.

You have all these ... you can hardly see them. All these gates, which prevent, or avoid people without the tickets, to go into the metro and the train stations. So it's all kinds of designs, which are explicitly representing some values, good or bad values, or desirable or undesirable values, but they are implementing those values in a very strong way, or they are nudging you in taking some kind of behavior. And in the same way as this very physical artifacts have policies, we have the same kind of policies, or the same kind of values built into AI systems. Thinks about moral dilemmas, and moral overloads, so how we build systems to deal with the fact that you cannot really solve and achieve, at the same time, all different types of values.

And issues about bias, about inclusion and accessibility, and so on. I'm sure you all have seen in the best, the moral machinery developed at MIT, which is trying to identify how self-driving cars should make decisions, or will make decisions based on what we, as people, think is desirable or not. Of course, when we develop algorithms to classify dogs ... and now you cannot really see nothing anymore. They will classify and identify dogs in all kinds of scenes, even in a plate of spaghetti. You cannot really see, there are all kinds of added, showing that it appears some dogs there.

So, excuse me. So we are doing a lot of design, and a lot decisions intensionally, or unintentionally, which really make a difference. So before we really look at the design in urban planning, is one of very concrete area, we should really look at what are the values and ethical requirements, which we are putting there. And are those values that we want to implement, or to enforce things, which it must be implement, should be implement or could be implemented. Also, if we look at really at ethics, and go back to philosophy, there are many different ethical systems which all lead to very many different decisions, so a Utilitarian AI system will take very many different decisions from a Kantian AI system, or Aristotelian AI system, because they look at the way of making decisions, the way of deciding what is good, what is morally acceptable, in a very different way.

So those things are implemented, desired or not, in our systems. So one of the issues there as well, is social acceptance of what we agree that it's okay. We vote for governments by mechanisms of social acceptance, and social choice, and that is not necessarily the same as what would be morally acceptable. And I've seen both in the Netherlands and US. We have seen recently very many cases in which that is not effective ... aligned.

So can AI systems do good? Do good in urban planning or other ways? It depends a lot on how we define, and how good is understanding those situations. There are cultural, individual, and situational difference, which should be taken into account. And the society for which we are designing is shaped by what we do, but it's also shaping our designs, and I think that's one of the ... I'm finishing, yes. One of the things is we can try to implement ethics, or moral reasoning, in our systems, but still just knowing about the ethics doesn't mean, necessarily, that those machines will be behaving ethically themselves, or representing, or implementing the ethical values that we think, or we would wish them to implement.

And like I already said, yesterday, responsible AI is not only about the systems that we develop, but also about the way we develop those systems. We are the ones who are responsible at the end. And yesterday, in several talks, I heard things like the algorithm told me, or the machine decide, or the outcome has determined that, which very much led to single out this sketch, which we probably all know. If you don't know, please go on YouTube and look at it, in which she keeps saying, computer says, "No." And what ever you ask her, she will say, "Computer says, No." It is very easy, and very ... trying to excuse yourself from the decision.

But I think we really need to look at AI from the perspective of accountability, responsibility, and transparency. And we, we are developing those systems, are the ones who are responsible for what the systems are going to do. Thank you. So before I continue, maybe now is number two speaker for the paper number two, Syed is here now? Yes okay. And Amin is here? Not yet, so we continue until number ...

Speaker 3: Number three.

Virginia Dignum: Number three, and then we see if we changed number four, and number five. Okay so Sujoy, yes.

Speaker 3: [inaudible 00:12:25]

Virginia Dignum: Okay.

Speaker 3: Let's thank Virginia Dignum.

Virginia Dignum: With the [inaudible 00:12:36] standard, is a suggestion. Maybe it will become a standard, but, so this the basis for European network of excellence, which is joining 27 countries in Europe, and outside Europe, there's few, and around 100 researchers are involved in that network. So this is kind of the cardinal message of the network, so we hopefully, after the project starts, it will become more visible.

Audience Member: The president of the United States, very interested in.

Virginia Dignum: We do have some United States participants as well, yep.

Audience Member: So on one of your slides you say we are responsible, so I guess the question is, who is we?

Virginia Dignum: Yeah.

Audience Member: Is it us, the designers? Is it the AI system itself? Is it the users? All of the above?

Virginia Dignum: Yep. All of the above. We as researchers are responsible, systems should be able to at least be able to indicate what they are responsible for, and what they not responsible for, so it should be clear what are the kind of scope of responsibility of a system, but as designers, as researchers, as developers, as implementers, as users, as owners, as manufacturers, I think we all have some responsibility in the decisions, and the ways these systems are going to affect our lives. And there is now all this huge discussion in the media about how AI is going to take over the world, and replace all us by robots, or whatever other kinds of very dystopic views of the future. And I think that the best answer we can have to that kind of discussion, is really to assume, and to indicate the responsibilities, and also the limitations of the responsibilities, at all levels of the design, yes.

Okay, so then we welcome the first speaker Sujoy Chatwterjee, on smart city planning with constrained crowd judgment analysis.

Sujoy C.: Thank you. So good morning everyone, part of my talk is on smart city planning with constrained crowd judgment analysis, and this work was done under the supervision of Dr. Anirban Mukhopadhyay, Associate professor of Department of Computer Science and Engineering at the University of Kalyani, and Dr. Malay Bhattarcharyya, assistant professor of Department of IT IIEST, Shibpur, India.

So these are the brief overview of my talk. So let's suppose that someone has posted a question, in social media, and he's asking that ... he wants to know, and find the answer of that is the question is, so showing my son the Disney Aladdin would not make him want to join ISIS, will it? So I'll give you the answer of this question, so we know that crowdsourcing can help us to understand, to solve the different real life tasks, we have the problem. So it can help us in harnessing the power of crowd and utilizing the [inaudible 00:15:47] and resources we can solve different type of real life problems.

So he can ... the person can outsource this problem, this question crowd worker, and he can optimize all the feedback from this. So now suppose that there are three possible options, there is yes, no, and skip, and now the question is there how to, if we have multiple opinions, then how to get the final answer, how to figure the final answer from this multiple opinions. So you see that according to the majority, voting no will be the final answer. But, so the judgment analysis is basically ideal to solve this problem to find out the [inaudible 00:16:26] judgment from multiple opinions. SO it's a general way of learning of [premonition 00:16:29] from multiple opinions.

So in this crowdsourcing environment, there is a response matrix that contains ... on the left side, the left hand side, you can see that there are so annotators, so that they are lying rows, and there are some questions, that are columns. And the self already knows the opinions of discussion for particular annotators. So there are five options. There is yes, no, skip, unsure, and I can't tell. And on the right hand side you can see there is a judgment matrix, that's where the rows there are some questions, and basically there are five options for this response matrix, so there are five options here.

The self already knows the weight of this particular option, for a particular question. So you can see that for question q1, what will be the final answer, because the what will be the maximum. Wehat is maximum for this question? So what will be the final answer, for the question q1? Now, in this problem, you can see that the response matrix have the options, yes, no, unsure, I can't tell, and skip. Now in different real-life problem, you can see that there are some questions that have some subparts that comprises some of the component ... subcomponent.

For example, suppose that a government organization is trying to install three ATM counters in your city, or in your locality. So he gathered the information about the appropriate location, or exact public demand. To know that exact public demand is not easy. So to seek the public opinion, exact public interest, over the ATM counters, you can outsource this problem to crowd. So that they can give their possibly options, possible locations. And there are some constraint. Again that, because two adjacent ATM counter cannot be side by side, so there should be some specific distance between any two ATM counters, so these are some constraint.

So although the [straight 00:18:31] of their approach deals with this problem, but in those problem, the question are basically of single or binary opinion, single or multiple opinions, but the question don't have any component. But in this constrained case, this is the [inaudible 00:18:48] This was the normal case, and this is the constrained case. In constrained case, for example, as we are discussing about three ATM counters, or five ATM counters, so there are three double ... this is called triplet of doubles. So because there are two possible ... three locations, with X and Y coordinate. That is for first location. This is for second locations, and this is for third location, and there should be some distance, and there should be some relations for doing this.

So while attempting the question, the [calculator 00:19:19] should also satisfy the constrained also. So this becomes a challenging problem. Now the problem is that how to indicate this solution, from this multiple solution, how to indicate the final solution. So in measuring response in the normal case, the response matrix, we can see that if we can apply majority voting for the past question, yes will be the final answer, because two percent are giving yes. But in this case, majority voting fail. Although the ... all of the straight up data approaches fails here, because if the number of range becomes high, then it will be too much tough to find out the majority voting, and it's a very less chance to repeat this opinion for all of these component.

So there is no chance to apply majority voting, so we have to rely on another method. We are proposing another method. So these are some of these challenges, and these are the problems from relation, and other thing is that, in this case, for this response matrix, for normal case, there are some options there also. But in this case we are just trying to set the range of this option [inaudible 00:20:27] There is a starting range and ending range, but we ... a starting point and ending point, but we don't know that what will be the final options. So first we have to help with the options, because without these options there, we can't hear the judgment matrix.

So in this problem formulation, there is a set of questions, a set of annotators, and we are the set of opinion vector, we have created this opinion vector by base end binning, to find the optimal number of options. So finally we have to compute the final [inaudible 00:21:00] judgment for each of these question. So we have to compute the opinion vector, we have computed the Gaussian binning, to scale this large scale of opinions into a small vectors. But due to the ...

This is the PGM model, we have considered annotated accuracy, and question difficulty, and these are the question component of different questions. And this is the judgment matrix for this [to be test 00:21:30] that is for X/Y coordinate. X coordinate this is row, column wise this is Y coordinate, but due to the time constraint, I am not going to the deeper end of this model. There are different factors, that is accuracy, and question difficulty. We have used the logistic formula, and we have used the expectation estimation algorithm here. And the accuracy plays a vital role, and for this case, we have used the coverage of this ... actually the coverage of this point, that is coverage of this ATM counter, for example, that if the points are well distributed, that means that it will be ... everyone will be served by these ATM counters, so that's what you call area of the entire, enclosed by this three point will be taken into account.

So for this case of creation of this dataset, we have posted a grid map, and this is the question that a top US university wishes to start the three extension center, and please give your opinion about the possible location, there should be some constraint. And after that the creation of this judgment matrix, we can get some score of this matrix based on the ... and we can to sort the score, based on this posterior distributional [inaudible 00:22:45] and for each of these component we can gauge some solutions, and the best solution is called a Rank-1 solution. And then we are computing this Rank-1 solution for both of these component. And we can check that whether this satisfy the constraint.

So for example, it will be Rank-1, and for this component two, it will be Rank-2, and for component three, it will be Rank-3. And if you satisfy that constraint then it will be treated as the final answer. So these are the opinions for [the faster 00:23:15] this is the [roaming 00:23:16] of options. That is, we are getting these options from multiple opinions, multiple crowd worker, this is after base end binning, and for in this case we getting the Rank-1 solution, and this is for Rank-1 and Rank-2 solutions. So for this case we have considered the accuracy, and the question difficulty of this ques- ... of accuracy and of the [inaudible 00:23:43] and question difficulty, but we can try to find out the vastness of the [inaudible 00:23:47] and other different parameter that can effect the overall judgment, and this is the references that I have used. And thank you.

Audience Member: Because you, in the end, get three locations chosen, but are they the best locations? Are they usable locations? [crosstalk 00:24:13] Do you know that people that actually vote for the things have any influence, maybe none of them has anything to do with university, but think it's funny to ...

Sujoy C.: For the first case there is ... when we are finding the accuracy of this annotator, first to the mean value is computed, and the deviation from mean value, to annotate that is our ... the quality of this annotator are actually computed, so this can remove the annotator who was ... or the crowd worker who are not giving their opinion, or who are not confident on their opinion. So in this way we can justify the quality of this work.

Virginia Dignum: How does your work compare to [counter set 00:25:00] mechanisms?

Sujoy C.: There is nothing about [counter set 00:25:06] mechanism.

Virginia Dignum: Because you are ranking, and then trying to be the creator.

Sujoy C.: Well ... no, no, no. Yeah, this is the-

Virginia Dignum: The average of the rankings.

Sujoy C.: No, we have not used [counter set 00:25:13] ranking. We have just sorted the posterior distribution [inaudible 00:25:17] for a particular option, that is for judgment matrix.

Virginia Dignum: Yeah.

Sujoy C.: For this case, if this is the maximum posterior value. 0.06. That means of 40 and 40. 40 comma 40, will be the faster options for the component. For [analytic 00:25:35] component, it will be that 40 comma 40. 30 comma 50 will be the best possible posterior distribution. So we are sorting this value in descending order, and the top will be the Rank-1 solution for this particular component. And in this way we are combing this ... checking the satisfying, whether the relation is satisfied or not.

Virginia Dignum: Okay. Any other comments?

Audience Member: Just a question about the possible solutions that you can get from the system. Is it true that the solution that you get, will match one particular annotator solution?

Sujoy C.: No.

Audience Member: Or will the event be a combination of-

Sujoy C.: Yeah, it's a good question. Because I was expecting this question, because if the opinions are diverse, then it is very less chance that opinions can belong to any of these particular [datically 00:26:38] match to the any of these particular annotator, and if the opinion center becomes very less, than there is no chance that it might be from this opinion vector, we are getting from this input.

Audience Member: But, there is no sense of centroid combina- ... that it's going to be a choice that someone has actually said, not an average of multiple choices.

Sujoy C.: No, for the first case ... for the initial ... 'cause as you saw it here, iteration wise, we are computing the accuracy. For the first case we are computing their deviation from this mean, after that, accuracy of this annotator, or the measure by the coverage. So after some iteration it can be seen that ... you can see that from this graph, from this plot, that it may be that this is not the exactly, for the red cross, this is not in the mean [case 00:27:30] but, this is left aligned, that means that although the mean solution exists, but is not a good annotator. That's why some other solutions are treated as the final solution, based on this coverage and accuracy.

Audience Member: Okay, thanks.

Virginia Dignum: Any other questions? Okay.

Audience Member: In know you're running out of time, but I wanted to know what's on the X and Y? What's on the X, and Y axis, and how to read these graphs?

Sujoy C.: Yeah, this is the X axis for location, X coordinate, this is for Y coordinate.

Virginia Dignum: So that map plot. That-

Audience Member: What are the different colors?

Sujoy C.: The different colors means there is a opinion ... There are three possible locations, that is for first rating encounter, the red. The green are second rating encounters. And the blue are these third rating encounters.

Audience Member: Thank you.

Sujoy C.: Thank you.

Syed Ali R.: So, so far we have talked a little bit about the GDPR regulations that EU has just released, and so, in this paper, what we wanted to talk a little bit about is what are the more relevant sections that we found in that regulation, and what are the things that we are concerned about, and what we think these regulations might inadvertently cause more discrimination in the machine learning applications that they're trying to mitigate. So let's start off first off by looking at ...

Okay, so while that is being sorted out. So what basically we want [inaudible 00:29:24] Ah, yeah. So what we first wanted to define was like what is discrimination? And there are three main components of what discrimination is. There is that the algorithm, whatever you're using, has a significant blind spot. The second thing is that, that blind spot adversely affects a particular sector of your population. So there are three things, it adversely affects people, a certain section of people, and it is arising out of a blind spot in the algorithm, and we are defining that as discrimination in a machine learning algorithm.

And so the GDPR regulations are actually trying prevent this sort of discrimination, and they are in that rim, they have come up with a very long, detailed update to regulations they introduced back in 2007, and these were ... Oh, brilliant, thanks.

So we have looked at this. Now in the GDPR regulations that they have introduced, there are a couple of definitions, and these two are the most important ones that I think we need to know, that concern processing, and profiling. So both these definitions have very broad ramifications, these are very general definitions, and I'll just give you a moment to just read them as they are.

So the first one, which is about processing, concerns that any set, which is whether or not, by automated means, is subject to this. And it includes things like collecting, recording, organizing, structuring, storage, adaptation, alteration, retrieval, consultation, use, disclosure. So there's a lot that it's covering, in processing.

And then in profiling, what it saying is that, any data subject, who is subject this processing, is profiled by an algorithm, when they're subject to any decision that is made by that processing, so that's profiling. And in particular, when we try to analyze, or predict aspects, about the person's performance at work, economic situation, so things like insurance checks, health, so any health applications. Personal preferences, so ads targeting, interests, reliability, behavior, performance, so all these things, including locations, or movements, so including people who might be migrants, and you want to track things. They're all subject to this regulation.

And so we can see it's a very broad regulation, and that's why it's important to know some of the more in depth things that it covers. So the first thing that is, it's a right of access, by the data subject. So what this creates is that the data subject can decide how they want their personal data to be accessed, or whether they want that data to part of some processing. And so that's the first and foremost right that this regulation introduces, control over your data.

The second thing it introduces is the right to object. You can object, and opt-out of your data being processed by any of the means that were described in the previous slide. And the third, and most important things that it introduces, is automated individual decision making, and a right to explanation, essentially. So this is the thing that creates a right to explanation. So the data subject has a right that if they subject to a decision, that they can ask what exactly is the explanation of what they are being subjected to.

And this is a tough one. So now let's bring this a step back, and just look for a moment at what this means for the machine learning life cycle. So this is the broad general machine learning life cycle where the outside blue is what the algorithm does in this cycle, and the inside is what the data does in this life cycle, and they're both entwined life cycles. They both go through a process of integration, where we get the data, and the algorithm, we make updates. We implement it, we process the data, and all of these sections are being effected, but two of them are being effected the most.

In the data life cycle, because of their right to opt-out, and the data life cycle gets effected at the separate part of data collection the most. And in the algorithm life cycle the last, third, kind of regulation that I showed, the algorithm life cycle gets effected at this stage of the algorithm selection. And so what do these two things end up creating?

So the first thing ends up creating a discrimination risk, because the origins of the data are sort of now not representative of the populations you're trying to effect. Now think of it through an example, if you are a health care kind of startup, and you're trying to implement an algorithm, your data might not contain certain type of people anymore, because they have decided to opt-out for whatever reasons that they might personally have, that maybe this data, if I release it, has some insurance costs, but because now that data is out of your access, your algorithm can never capture those things, because it just doesn't see them, or enough of them.

So that's one thing that, when you have bias coming in, you will have bias coming out. And so inadvertently we see, and we think, that this regulation could create this inadvertent consequence. The other thing that it ... So, and how do you manage that bias? Well, you have to be very mindful of where your data is being generated from, and how are you using other datasets to fill those gaps? Because they're extremely good ... there's extremely good research happening right now, which is trying cover how to cover gaps where your data in insufficient, but then we need to be hyper-aware of the origins of the data.

The second thing that it creates, the black-box risk, it exacerbates that, because currently we are using algorithms, that we are making them more, and more advanced. More complex. We may not have interpret ability precisely for them. So maybe, for some applications, choosing a simpler algorithm, which is easier to explain, but, which might lack the nuance for the edge cases for which we created the complex algorithms in the first place. So there is a selection bias problem, and a lot of this work is sort of trying to start a conversation on what does this actually mean for us, as data scientists?

And so I thank you for your attention.

Virginia Dignum: Any questions? Yes?

Audience Member: Have you done any application?

Syed Ali R.: So, the applications that we ... my personal PhD research is sort of apart from this, I work with financial data, but over the last summer, in Chicago, I worked with Data Science for Social Good, and there, a lot of applications that we were working on, would have direct consequences if this regulation was, let's say, applied to the US. Because there we are working with school kids data, policing data, and if you just for a moment imagine that you take out whole chunks of the population, in that. Because if you have that right, you probably would want to exercise these rights, and it would create biases in the data, so that's where the, kind of, idea started, that these applications would most immediately get effected, by these regulations.

Audience Member: So there's a recent ... so recent work by Jon Kleinberg shows that some possibility theorems that, if you look at the criteria for what ... for, example seven [inaudible 00:37:43] guidelines fitting nondiscriminatory, you cannot satisfy all the criteria simultaneously. So I was just wondering, if getting into thinking that whether these regulations might actually be impossible.

Syed Ali R.: So that's an excellent point, and I agree with you, to a certain extent. That, maybe, there are a set of constraints, which might be impossible to sufficiently exercise, and one of the biggest sections, of the regulations, that might be very hard simplify, is the interpret ability problem, because ... and someone used this example yesterday, that going to a doctor, we kind of get a watered down explanation, but a lot of us leave feeling very unsatisfied, but that only happens because we don't have another choice. There is no alternative doctors, that exist, that we can go to.

But in the case of automated processing, there is an alternative, you just opt out of being processed that way. But that brings us back to square one, where then we are back into the realm of manual processing, and then the kind of progress that we are making, in data science, becomes, sort of, nullified.

Virginia Dignum: Just one issue on the bias, human generated data, or human related data, is necessarily, and by definition, bias. 'cause people are bias, and the bias appear in the data, that you collect, or that you store, or whatever. As you go about removing the bias from the data, you are kind of introducing the bias of the ones who are determine which bias should go out of the data. You understand what I mean? How do you plan, or what are the ways, to go about that?

Syed Ali R.: So there are three things I would say that relate to your question. The first thing is that machine learning, or data science as we know, is no magic. It won't automatically marginalize out human bias, it will propagate it, and it will formalize it, if there is that, and that is your application, then talking to experts is, of course, the way to go, and it can be very challenging at times, because you have a personal moral compass, and when we saw the predictive policing application, the sense that there is something icky about it, that arises, and you might have a mismatch with your client's expectation, versus what you want the world to be. And that's one big problem, which I don't have an answer for.

But the second thing is that the inherent data bias, of course there needs to be quality mechanisms that we, as data scientists, need to come up with, that this data now passes a quality threshold, and it won't have unnecessary biases, just that are artifacts of the data collection process. That's one thing. And the second thing that is very important, is to actually see, what does our error margin mean now? Because it's very easy to say, "Oh well, I'm going from a 95% to 95.1% error, maybe it's not that much." But like we saw yesterday, going from a 93% to a 47% percent per- ... or sorry, a 83% precision, is something that happens, but what is that 2%, that is maybe 20 peoples lives, so am I willing to make that call?

So now our precision, in data science, takes a whole new meaning, because you are no longer just concerned with 3,000 more images being classified, you're actually talking about 20 people maybe being released, or not released, based on your algorithm choice.

Audience Member: So I have a thought that emerges from this, that I haven't heard expressed yet, and from the perspective of social work, and other social justice oriented social scientists, we are not only concerned about bias coming out of these processes, but we're also interested in uncovering bias that may not have been known. So for example, in health care access, or in criminal justice, we have policies that are laws, that are intended to prevent bias. You know, you're not supposed to have policing practices, which are in fact racist, or sentencing practices, which are I fact racist, and so I wonder if one of the unintended good consequences, of this sort of data, sort of look that you all are doing, is that you may be uncovering biases in systems that think that they do not have bias.

So rather than, sort of, despairing that, oh, our algorithm, it turns out, in predictive policing, is telling us basically that if you're black, you're never going to get a ... you're never going to be released on your own recognizance. In fact, what that reveals, is that the system that was built off of the data of these practices, means that if you're black, you never get released on your own recognizance, and it actually quantifies, and measures, this sort of bias, that people don't think is there. And it strikes me that this is actually a whole new area for pursuit, that you'd ac- ... the pursuit of bias, in systems, and then obviously you want to try to then make decisions rules that come out of that, that don't have those biases.

But I think you all are in a position to identify things, which systems, which think that they are not doing, you know that they're not biased, are actually biased. And I'm wondering if this is, in any way shape or form, part of the dialogue.

Syed Ali R.: So I completely agree with you, and it is, especially a particular paper that really inspired us, and we reference it, it's called: Is tech racist? The fight against digital discrimination. And in it, or the work behind it, tries to uncover whether fundamentally the process we are trying to describe is inherently discriminatory, and so are we going to just ossify the discrimination? But I'm aware of time, so thank you so very much for listening.

Daniela Rosu: Let's see. Actually yeah, I can use yours. Okay. All right, let's see. I don't like how ... it like erases.

Hello everyone, good morning. My name is Daniela Rosu, I am the first author on this paper. I'm with ... I'm a [professor 00:44:44] fellow with the Center for Social Services Engineering at the University of Toronto. We're a group of computer scientists, mostly, but also industrial engineers looking at applying computer science, and engineering principles to efficient, and hopefully effective, design and delivery, of social services. It's a very ambitious goal. We have some people from the faculty of social work, as well, working with us, and what motivated this particular work.

So just to let you know, this is a big project that I'm going to talk about today. The individual technical parts are actually being submitted as technical articles at various conferences. This is just an overview of the problem we're trying to tackle, and the solutions that we're proposing. What motivated our work was the fact that an increasing number of people, we call those vulnerable populations, are in need of products, and services, they can no longer afford to buy. And by vulnerable populations we really mean, you know, the people who are not truly privileged, they live in poverty, the frail elderly, people with very complex health needs, the homeless.

And for us, in Canada, it's a real issue ... well, for Europeans as well. The fact that last year, according to the estimates of the United Nations, 65 million people, more than at any other time since the Second World War, so 65 million people, were refugees. And Canada got about 40,000 of them, in a very, very short span of time, most of them actually settled in Toronto, and many of them were privately sponsored, which means Group of Five, Canadian citizens, would sponsor a family to come into the country, and they had to provide everything for them, car, apartment, furniture, clothing. Get them enrolled in school. Get 'em tutors for the English language. Everything you can imagine had to be supplied by this five private individuals.

But on the other hand, we live in a very rich society, at least here in North America, and also in Europe, we're very privileged. However, we're not very good at distributing this wealth of resources to the people who actually need them. So, you know, we figure that since we bear it ... like our center is a center for social services engineering, we'll try and look into this issue, and we figured that the things that we're really, really passionate about solving are the following.

How do ... not just determining the active demand, what people know they need, what people don't really know they need, yet. And an example that my [professorial 00:47:32] supervisor is very fond of giving is, you have this Syrian pregnant lady coming with her family into Canada, she doesn't know she needs a car seat. She doesn't need ... she will need a vaccination for her newborn. She doesn't need a number of things ... she doesn't know she will need a number of things. So this is latent demand that we're trying to actually anticipate, and fulfill.

So we want to be able to identify the problems that people need to have solved, not just how we're going to deliver to them, you know, a mattress, or a tutor for their kid, or something like that. Also, how do we discover the whole supply side. Again, we do know what the active supply side is, but how can we learn models of what people might have. Who are these people who might have goods, or be able to provide services? And they don't even know it, right, because nobody has asked them, they never actually thought about it.

And of course, after we figure this out, how do we best match, and allocate them? This is very important. How do we allocate supply to demand? So we have a plethora of demand and supply, but choosing the best match, based on needs, based on subsumption, or equivalence of product, services functionality, et cetera. This is very, very intriguing, and I'm very passionate about this, you can tell. And of course, last but not least, how do we transfer the product, and services, from where the supply is to where the demand is? So a lot on our plate.

While we want to actually propose a solution, is the next generation marketplace. We're not trying to be Kijiji, we're not trying to be a Craigslist. There are problems that Craigslist, and Kijiji solve very well, this is not what we're all about. So we want to be able to match the demands side, and the supply side, and whenever needed, we want to be able to schedule volunteers for either the delivery of goods, or a provision of services. Of course this is all very ambitious, we're years far from being able to solve that, but we started looking at some concrete examples that we can actually tackle.

One of the first ones was ... no, it started from the foundation, foundation work. How do we represent demand and supply? And here we have a number of issues, because in order to actually attempt high quality match of supply and demand, we really need very detailed descriptions of what the goods and services are that we're trying to match. And here, if you've ever looked on Craigslist, or Kijiji, or even on the websites of service providers, like government organizations, or NGO's, the terminology that people use is very ambiguous. Even, we humans, scratch our heads sometimes and are like, what is this? Not sure.

Many of the terms are really adhuc, like you ... it's in very contextual, like you wouldn't even know from one context to the other what they actually mean, and how they relate to each other. In terms of harmonization, how do we integrate all those descriptions, and make them work with each other? And of course because many of us are in knowledge representation and reasoning field, classical AI that is, non-machine learning, well we're doing machine learning as well. We use formal representation, or anthologies, for goods and services, and not just that, but this is something that we're actually focusing on right now.

As far as ... so this is sort of a progress report, this is where we're at. I told you the lofty goals that we have in mind, I'm gonna tell you now a little bit about what we're actually doing. In terms of matching and allocation, the biggest issue that we have, is that there is no perfect match between supply and demand, as far as we could tell, which means we have to really figure out a way to define, first of all, what the best match is, and that is, again, contextual. And also because we wanted to make this as unbiased a process as possible, we figured we would try and use formal methods for that, so we are trying to use the equivalence and subsumption of product and services, as given by formal representation, so [versatile 00:52:02] logic, in this case, for us.

We also ... maybe I ... yeah, no. This we can skip. So as far as representing and solving our logistics problems, so how we have this fleet of volunteers, they have cars when you need properties, they also their new constraints. What is it that we're trying to solve? So even representing this problem, actually had it's challenges, and I'll talk about it a little bit later. So, the very basic would be, now how do we schedule the pickup and delivery of goods? And you'll see that it's actually not that easy. Our colleagues, Professor Beck, and his student, use constrain programming for doing so. Actually it's a hybrid model, I'm gonna talk about it a little bit later.

Coming back to actually my area of expertise, we tried to ... we started digging deeper into issues related to representing products and services, so we wanted to know. We divided the pie into smaller chunks that we could ingest, and this is like the smallest one that I could actually look at first. Some of the issues are actually quite intriguing. The demand sides need to know the condition of the used goods offered. Again if you look at Craigslist and Kijiji, people describe the different goods that they have to offer, in a variety of ways, very creative ways, because many times they want to actually ... they want to be truthful and say, yeah there's a little bit of damage here and there, but the language that they use is not ... they make it sound like the damage is actually lighter than it is, and especially if ...

So one thing that we discovered is furniture banks, in Toronto, do not accept baby furniture of any kind, so you cannot find there a cribs, you cannot find strollers, you cannot find car seats. We have to source those, for the refugees, from the general population. But we need to know, right? Was this a subject ... was this car seat subject of a recall? Is it in good condition? Is it something that people can use? Same with the crib, is it safe? And it's just based on the descriptions people give, that's not immediately obvious how to do.

Also, checking eligibility for transportation. So recall we have a fleet of volunteers, and some of them have pickup trucks, some of them have minivans, some of them have just, you know, sedans. For us, one of the issues was, okay, a lot of people, at least in Toronto, use Ikea furniture. Can we dismantle those pieces of furniture, so that they actually fit in the car? So that was something important. This is not information that people usually give you on Craigslist, on Kijiji. We needed to be able to source that. And also when we schedule people to be transported, let's say to medical appointments, if they have mobility issues, we need to be able to do that.

How many minutes do I have?

Virginia Dignum: No.

Daniela Rosu: No, okay. All right so, just one more. So we've done as much as we could, so we have a prototype, we implemented that using a Blackboard architecture, and it's online. We're actually piloting this with an NGO in Toronto, and we want to do some future work as well. Thank you. Wow, that was fast.

Virginia Dignum: Okay, questions?

Daniela Rosu: Actually, maybe this one, yeah.

Audience Member: Hi, first off I would like to commend our neighbors to the north in their approach, and proactive response to the Syrian refugee crisis.

Daniela Rosu: Thank you.

Audience Member: It's been much more effective, I would think, than our country has been able to do, so ... Second, there was a very good New York Times article by Cantor and Einhorn, this past Saturday, about what's been dubbed in Canada, Month 13, which is these year long adopted refugees.

Daniela Rosu: Yeah.

Audience Member: Are ... have a contract of one year, and the transition from month 12 to 13 has been a social, kind of, crisis, at least in Canada, about what to do about this. And now these people are here, they've been here, they've been assimilated, to some degree, the culture, and their society, how do you deal with that? Is there any role in this marketplace, since you're dealing with economics of refugee management, is there any role that that can play to help mitigate that problem?

Daniela Rosu: Absolutely. So this work that I've just talked to you about, we actually have programmers with us. So at the beginning it was a research prototype, that I and undergrad students, helped develop. Now we have a full time programmers working with us, so this is being piloted by a few settlement agencies in Toronto, and that's exactly to try and alleviate this shock, 'cause after 12 ... so what I didn't say, is that, this five individuals, who are private sponsors, have the obligation to take care of this family for 12 months. Everything this family needs, the sponsors have to provide, but their legal obligation stops after 12 months.

So ... but, these people still need a variety of services and goods after that. Yeah, absolutely. And this is what we're trying to do. It's an ongoing process, it's not like, yeah, we got these refugees in, and now, you know, it's not our problem anymore.

Virginia Dignum: Any other questions? Do you have any results already, from the work you are doing with the NGO?

Daniela Rosu: No, not yet, 'cause it's like literally, so I was ... I just got a text, as I was sitting here, and somebody was telling me, yeah, we have the go ahead, we'll deploy this. So, not yet. And we're really looking forward to having this use, because all this work, which requires machine learning, cannot be done without the data that we ...

Virginia Dignum: Yeah.

Daniela Rosu: 'cause right now, we just have a run, and test, just with synthetic data.

Virginia Dignum: Sure.

Daniela Rosu: So we're doing very good, but synthetic data. Yes? I think it was a ... oh, sorry.

Audience Member: So it's wonderful that you have this center. I guess, as you know, there are these centers that are coming up, we have a one at USC. Is there something that can be done 'cross centers? Is there some, like you know-

Daniela Rosu: I would love to.

Audience Member: Some sort of a collaboration, but I'm not exactly sure what structure it would take, and it certainly, you know the dean signing a memorandum of understanding, or something, that doesn't really kind of advance scientific collaboration, or other kind of collaborations, exactly, so it sort of got to be researcher to researcher, but I don't know ... so I mean this might be something that could be taken up tomorrow, but-

Daniela Rosu: Oh, absolutely. I would love to talk to you, and also to their dean about it.

Virginia Dignum: Okay, let's go. Yeah, okay. So she start, and then it's him and then ...

Audience Member: I just wanted to know how you were getting the additional metadata that you need to match supply and demand, in this context?

Daniela Rosu: So we've had a lot of interviews, we actually talked to settlement agencies, and in particular with settlement workers, 'cause they are the ones who know what the needs of the refugees really are. I haven't personally talked to refugees directly, but I also talked to some of the private sponsors. Actually, one of the web developers that we have was part of one these Group of Five, people, so he knows intimately what kind of struggles they had, in order to provide for those people. What they needed.

And also, in terms of ... Actually, what I didn't get to talk about was representation of services, 'cause we want to also bring AI planning into the fold, 'cause some of these people have really complex needs, in terms of services, and if you think about the ultimate goal, is to integrate them in to Canadian society, then you start [decomposing 01:00:07] that into actually practical tasks that you can achieve, and then you can see how this could actually, you know, it kind of lends itself nicely to the planning.

So for coming back to services, we again looked at what the providers themselves, NGO's, or the government, describe as having services. And, you know, eligibility and [inaudible 01:00:29] care and all that, gets enough. 'cause we were looking at service composition.

Virginia Dignum: Questions, stick to questions.

Daniela Rosu: Sorry.

Virginia Dignum: Can you?

Daniela Rosu: I didn't expect [crosstalk 01:00:36] so I'm happy to see.

Audience Member: How do you deal with the all kind of practicalities? So there's a lot supply for vets, and not enough demand for it, at this particular moment, do you have some storage room, is there a-

Daniela Rosu: Yes, so-

Audience Member: Do you have this contact with NGO's?

Daniela Rosu: Right, right. What I didn't get to talk about is the fact two, one of the things we're trying to achieve with this, this is really the gift that keeps on giving, 'cause we're extending way beyond our capacity. We're actually thinking of helping NGO's, that work on similar problems, become a virtual organization, so seamless integration. And for this particular issue, that you raise, we approached furniture banks in Toronto with the idea of building a virtual infrastructure for them, so exactly. The extra capacity that we cannot process right now, we could actually divert the warehouses.

Virginia Dignum: Let's move to the next question.

Audience Member: Hi, I was just wondering, you mentioned that you were focused on refugees that were settled by private donors.

Daniela Rosu: Yeah.

Audience Member: As far as I'm aware, the settlement organizations also ... they generally settle refugees by using governmentally sponsored. 'cause I volunteered at one in Ottawa.

Daniela Rosu: Yeah.

Audience Member: And so I was wondering whether this will be also available to them, for them to be able to-

Daniela Rosu: Absolutely.

Audience Member: Great, okay.

Daniela Rosu: Absolutely. Actually ... well every NGO in Canada is really, mostly, funded by the government.

Audience Member: Yeah, exactly so-

Virginia Dignum: Okay, so-

Daniela Rosu: What I didn't tell you is that we have 60,000 NGO's in Ontario alone, which is-

Virginia Dignum: So, then we continue at the coffee break with ts . [crosstalk 01:02:16]

Daniela Rosu: I'll be happy to talk with anybody who has questions. Thank you.

Virginia Dignum: Thank you.

Speaker 8: You can either hold it, or leave it on there.

Amin Ghafouri: Sure.

Speaker 3: Did you press J-U-N?

Amin Ghafouri: Yes.

Speaker 10: [inaudible 01:03:07]

Amin Ghafouri: We got it. Uh, the colors are.

Speaker 10: I think that's with the [inaudible 01:03:24] projector.

Speaker 8: Yeah.

Virginia Dignum: It's wrong color.

Amin Ghafouri: Oh, okay. I hope this doesn't cause any troubles.

Hello everyone, this is a work by me, and my colleagues, at Vanderbilt University. In a smart cities, realtime traffic sensors are used for different applications, such as realtime control of traffic signals, or route planning, however, traffic sensors are prone to failures, and failures can result in degraded performance. The challenge that we face is to quickly, and accurately, detect faults, and anomalies, in sensors.

However, detectors are imperfect, that is there are detection errors, which are false positive, and false negatives. As you might know, a false positive ... excuse me.

Speaker 8: Sorry, could you just use the handheld?

Amin Ghafouri: Sure, sure. [inaudible 01:04:30] just gonna set my timer, okay.

A false positive means raising an alarm when the behavior is normal. On the other hand, a false negative means raising no alarm when the behavior is anomalous.

Speaker 8: Sorry could you move-

Amin Ghafouri: Oh, sure, sure I forgot, that was a good point. I'm not used to this.

Speaker 8: I know.

Amin Ghafouri: [Show of hands 01:04:52] While it is desirable to reduce both false positive, and false negative ... I got it.

While it is desirable to reduce both of them, it is shown by the security literature, that there is a trade off between the two, that is, decreasing one, increases the other. So the problem that we are going to solve is minimizing losses due to false positive, and false negative errors, considering smart city application, in this work we consider route planning.

So given a set of queries in a route planning application, the goal is to find the shortest paths. Let's say Google Maps, or Apple Maps. And a fault is ... oh, the figures are really changed. And a fault is modeled, as you can see, so the measured value if it's M, is different from the actual value A, because of the fault epsilon. And this results in the routes that are computed being suboptimal, you know, which causes increased delays. To be able to detect anomalies, we introduce a detector. Our detector has two parts, a predictor and a statistical test.

For the predictor, we make the assumption that the number of sensors that are anomalous is low, which is true in practice. So then we can use the value of nearby sensors to do the prediction, because the nearby sensors give us a relatively correct measurement. We use Gaussian Process regression with ARD squared exponential for the prediction. So ... okay ... to summarize, we use the measurements of other sensors to do the prediction, and then we compute the differences between predictions and measurement to compute residuals, and then we feed these residuals, these differences, to a statistical test. We use CUSUM, which is well known, and is defined as ... it introduces a variable S of K, that uses the previous value of S plus this residual to compute that the current value of S, and if it's greater than a threshold, we raise an alarm. We raise a detection alarm.

So this kind of shows the whole idea, that the overview, for a detector with threshold error of K, so the threshold can change in time. An alarm is raised if S, if there's a statistic S, is greater than threshold. So if the threshold is small, our detectors is highly sensitive. If it's large, it's not sensitive, which means it results in too many false negative. So our goal is to find that the right value of threshold, the sweet spot on this curve, this trade off, that minimizes the loss, due to false positives and false negative, considering the application, in this case route planning. So we change the threshold dynamically, at every time step, to minimalize the losses.

So to do so, we formulate what this loss mean, so in the case of route planning, it means extra delay in travel time, so we formulate cost of false positive, and cost of false negative, meaning when there is a false positive, we discard the measurement, we use the prediction. If we are doing the right thing, what is the extra delay? That, that's what these equations mean, [inaudible 01:08:23]. And once we characterize the cost of false positive and false negative, we define the loss. The loss, L, that given a threshold, instead of route planning queries, it just says how much potential extra delay is caused.

So we compute this loss, the first term is for loss due to false positive, the second is loss due to false negative, and we define optimal threshold as the threshold that minimizes the loss, minimizes its expected delay. So since trade off curves are non convex, this problem is challenging, but we propose an algorithm that preforms pretty well. It takes as input, the queries, the predictions, and the measurements, and it computes the error of the thresholds, and it feeds it to the algorithm ... to the application, and it does it at each time step.

So we are also able to ... using this framework, we are also able to identify critical sensors, which are sensors that cause high losses in our system, and identifying them lets us know that these are the ones that need to be secured, and be made robust, first.

So we talked about the theoretical framework, now let's talk about the case study, what it means in real world, in real life. We use a real life datasets of downtown LA, home of the conference organizers. We use 115,000 data points, is traffic data for two weeks, first week is for training, and second is used for testing. We also consider two fault models, one as three to seven percent error, the other subtracts seven to 13 percent error.

This is our detectors performance. This is similar to ROC curve, but instead of just true positive, we false negative. And we can see that our detectors perform very well. Now our next goal is to ... so each point on this curve corresponds to a threshold, right? Our goal is to find the right threshold, the right ... the sweet spot, that minimizes our losses. So we consider a scenario. First, we use the real life, an opens source route planner, that is actual use, you can use OpenTripPlanner, and we consider 1,000 route planning queries in a day, and fixing this scenario, now we run our algorithm at each time step.

And at one time step we see the result, and we see that our algorithm is able to find the optimal threshold values, that minimize the loss. So it is expected, you know, 20 seconds, and our algorithm, in this case, performed the global optimum, and we also identify critical sensors, so this sensor is critical, and we see that our dynamic approach, meaning changing threshold at every time step, performs much better than fixing the threshold, which is what is commonly done. You fix the threshold, a decision threshold, and you're done.

We see that our approach reduces the loss by up to 40%. So this is the conclusion that we made, that our approach minimizes the loss by 40%. How much time do I have?

Speaker 8: You're almost done.

Amin Ghafouri: Good. So I'm gonna spend like 10 to 20 seconds talking about ongoing work, the extension to this work. So the application is really, it has nothing to do with the framework itself, it could be any application. As an example, we can apply this to traffic signal control. I really like this figure by the way. So we have a traffic signal, we have a traffic signal controller, and we have traffic flow measurements, and a predictor. And then there's this detector, that it switches. It says whether the measurements inaccurate, or the predictor.

So this is just an example, that choosing this threshold, effects what's going on here. If we extend this even further, it could be any system, any physical system, that you have a detector, a controller, and the application. And this framework, the idea is to minimize the losses due to false positive, and false negative, by selecting this threshold that affects everything. This is the ongoing work, the loss function would look like this. So here I relax the assumption on nearby sensors not being faulty. I haven't solved it, I have four to five minutes to defend my PhD, and good luck to me on solving this problem.

So in conclusion, we designed an effective detector using Gaussian Process, that performs well. We find that we presented an approach for computing optimal thresholds that minimize losses, due to false positive and false negative. We also characterize critical sensors whose failure has high impact on the application, and we implemented an evaluator, evaluated our approach using a real life, real world dataset, and a real route planning application.

Thank you very much, and please let me know if you have any questions.

Audience Member: So have you done any evaluation, in comparison to other techniques?

Amin Ghafouri: Good question.

Audience Member: How it is going, in respect to what other people are doing?

Amin Ghafouri: Right. So typically, in the literature, they consider static threshold, meaning that they just fix the threshold, and they just keep it fixed for the entire lifetime. So in this case, this static, means the best scenario, in a fixed threshold case. So the comparison is made right here. On left hand side, this is what we obtained, and on the right hand side, you see what is commonly done in the literature. And we see that our numbers are up to 40% better than the existing work.

Virginia Dignum: Okay, any other questions?

Is your work only applicable for fixed sensors, or could you also use it when you have moving sensors, like in the cars, and use the cars as the sensors for all the city?

Amin Ghafouri: Oh, okay. So you mean the sensors-

Virginia Dignum: The sensors are not anymore stuck in a place, in a cross route, but they are moving.

Amin Ghafouri: Of course, of course. As long as we can have a predictor for that sensor, we can implement a predictor that predicts the future.

Virginia Dignum: The neighbor sensors are changing all the time?

Amin Ghafouri: It's fine.

Virginia Dignum: Okay.

Amin Ghafouri: Yes.

Virginia Dignum: Okay.

Amin Ghafouri: So the predictor, also, it could be anything, so as long as you give me a predictor, we can do that.

Virginia Dignum: Okay.

Amin Ghafouri: But, of course, for a moving sensor, designing a predictor is a challenge itself. So perhaps I need to do another PhD to just work on ...

Virginia Dignum: Okay. So thank you and ... thank you again.

Karen Judd S.: Do you have a [inaudible 01:15:33].

Okay, so I guess I'm on. I'm not sure that I'm in urban planning, but I'm here. So, thank you very much. Anyway, Virginia said that she began her session by saying we need to take a step back, well I want to take a step even, perhaps, further back. Because when you're looking at the United Nations, you kind of ... the assumption is, actually, that you're looking at a global sphere, you know, a global picture, but I do want to note that it's perhaps not quite all that it seems. But I want to just also explain to you a little bit, because I do come from sure a different background, I'm not looking at the granule levels of all of the algorithms, and mathematics, although once upon a time I looked at that stuff.

My bubble's been a little bit different, so physics, and history, and philosophy, of science a few years ago. But since then I've been in the international NGO arena, and 20 of those have been working at the United Nations, I have my own version of an algorithm, for advocacy at the United Nations. I've often focused on change leadership, both at the personal level, and at the organizational level, because the two are interconnected, and I'd have to say that one of the things that is often understood, and yet somehow forgotten, in the process of engaging, and change, is that there's the personal level, then there's the organizational levels of change.

And organizational change is often much more challenging, and often ... we are all in some kind of organization, but we see through that filter often. But we often forget, also, that what we need to engage in is upgrading, or adapting, that organization, its structures and programs, as well as our own at the individual level, or team levels. Anyway, increasingly I think it's pretty ... none of this is rocket science, but I want to give it to you as a background to where we're going.

Where my main point is very, very simple, and I think it's kind of ... right at the top it says, it's a missing conversation. I think the AI community, the tech community, needs to be much more directly engaged with the global conversations that take place at the United Nations, but that's not always a simple process. I liken it ... I spent a number of years on the water, and I liken it to trying to transfer somebody from a small vessel, onto an ocean liner, or an aircraft carrier, not so simple. Nice idea, it could but safe up there, but to get from the tiny little vehicle up to the major one, is actually quite a challenging process.

But anyway, we've seen this, everybody can see this, it's an experience that pretty much everyone is going through. Technology, and its evolution, is allowing and enabling all kinds of new activities, and possibilities, and we're hearing that all the time in everyone of these sessions. But our social structures, and our legal systems, are kind of slow at changing, and in a sense, rightly so. But the fact is we've got this huge area where it's more like a gray, I don't know whether it's 50 shades of gray, but it's ... might be more.

But it's a large area where we still are operating, but there are no absolutely strict guidelines for that, so we rely on our own standards, our own ethics, our own morals, our own organizations. And to some extent we rely on ...

Speaker 8: [inaudible 01:19:40]

Karen Judd S.: What's the number? Five.

Speaker 8: Five.

Karen Judd S.: Great. We rely on all those kinds of things, for helping us figure out how to get through this area, and that's part of what we're wrestling with here, is how do we do that? How do we rely on? What do we rely on AI for, and what are the ethics of AI, and how does that get integrated into our systems? And how do we then even convince those who do not know anything about AI, and are perhaps suspicious of it, that maybe it is good, or maybe it's not good? We don't know. But there's this just huge massive area that we're needing to deal with.

So one of the things, and I'm very quickly going to go on into ... this is, who do we trust? Just some interesting little tidbits. Fear, motivated about 67% there. 45 motivated ... you know it .... The thing is, one of the things that's going on, when we're in this area of uncertainty, is we do, as human beings, and this is where I want to get into a kind of ... my little view of change leadership, and some of the dimensions that I think do inform our concept of the social good, a little bit. And I'd like to just quickly go over them. And that's difference, our drivers, and thew domain, or scope.

Again, not rocket science, just a way to help frame our sense of what the social good is. So the drivers acknowledges that we do have three brains. Evolutionarily we have three brains. So in the previous section, what we were looking at is how fear is driving a lot of our social reality. And that is very simply because we do have three brains. We are human beings with a reptilian complex, we do have our paleo-mammalian, and our neo-mammalian complexes. And that's wonderful, but we cannot ignore the reality of those underlying driving forces that are, in a sense, pre-rational.

So anyway, difference, just very simply, is looking at historically in larger organizations, in corporations, a lot of the management. Sustaining the organization, moving forward incrementally in a controllable way, is you know probably, the differences are not big the here. At the moment a lot of the differences that you're engaged in, compared to what the majority of society, are innovative, they're a big difference, and that big difference creates vectors, so to speak, in society.

So then you've got the domain or scope, so you've got individual level, you know families or tribes, depending on where you are. You've got you're community level, you've got your national level, and you've got your global level. And our organizations need to be structurally appropriate to the levels at which we intend to have that impact. But sometimes, in that structural appropriateness, it's not exactly ... we don't always tell ourselves exactly the way it is, and we've got ...

Speaker 8: Two minutes.

Karen Judd S.: Two minutes. So, okay. So you have to read this. Okay, so then back to this. And I just want to say that this is an essentially part of us, it helps us understand ... When I went back to Australia, throughout the election process, of course everybody wanted to know what was happening in American politics, and first out it was, "Is it possible he could?" Then, "You mean he's really gonna run?" To, "You think he can win?" There was a ... and I had to give a very short explanation, and my explanation was, he speaks from his lizard brain, to the lizard brains. That doesn't sound so good, does it?

But, anyway, the point is there, we do have ... this part is us too, each one of us. So it's not bad, it's just that when we don't balance it out, and include the other components, and bring those into whatever it might be that we're doing. Okay, so anyway, real world problems, I'm so glad that we're solving them, engaging 'em. View it as very complex. I can't go into all of the details about how the UN works, but when I was meeting with the under secretary general of IT, just two weeks ago I discovered that there is a new initiative that she has begun, and it's called the Digital Blue Helmets, and this is where it comes into what's possible for the AI, and the tech community, as a whole.

Again, going back to that big ship, one of the ways to get from what you are doing, the granule level, the equations, the studies, the research that's going on, to the application that is useful for our world, is not such a simple connecting point. But the Digital Blue Helmet platform, that is newly established, allows for that. The caveat being, of course, the UN's ... anybody? Do you know UN structure? How it works? Okay, no.

Audience Member: No.

Karen Judd S.: Okay, believe me. Yeah, like everybody's going to believe me now. But anyway. Okay, so the secretariat, of what she is part, is like the big secretary to the member states, the general assembly, and the security council, and the economic and social councils, where the governmental representatives do their work. So the secretariat often has the largest global vision of the world, because they're attending all 193 nations, they're not just looking at the individual national interests of one nation. So they tend to have a global view. But that's not their job to do that. Their job is to service the member states.

So in establishing the Digital Blue Helmets, this is primarily to service the needs of the United Nations member states, luckily I've been so glad to see that everybody's been ... oh, putting that up there. And these are actually a way that the platform can be used, so that should a crisis arise, should assessments need to be made in any of these areas, there is the potential, as we go down the road, to do so, utilizing the Digital Blue Helmet platform, where you can get in there and do some of your granule work on specific issues. I just want to put that out there so that you are aware of it. It's only just beginning, and so it's not up and running per se, and there will need to be steps in between, because of this secretariat servicing the member states first, kind of concept.

But the point is, that there is now a platform. A way to get from our little boats, up onto the big ship, and if anybody wants to talk to me more, or has questions, thank you very much.

Virginia Dignum: It's a great topic for a coffee break, but if there are a few short questions before we go for a break?

Karen Judd S.: Yes.

Speaker 3: So is it the case, I mean I've talked to some people that work at the UN, and they usually tell me that, I mean, people there are not they're resilient to using AI, because they believe that it's not ... it's good where it is being used, and it is not as useful when you look at country level matters, and maybe they're right. I mean maybe when we talk of machine learning systems, I mean a false positive, as Syed was just mentioning, it takes a completely different dimension in when you're dealing with country negotiations.

Karen Judd S.: Right.

Speaker 3: What does the false positive mean? And so may their systems are not ready yet, and what do you think would be a good way to convince them, the people at the UN, for them to be able to use AI based systems, more, and more? Because right now I think, the people who are the leadership they are not from the age when AI was built.

Karen Judd S.: Right.

Speaker 3: Like they're from an earlier age, yeah.

Karen Judd S.: There's probably three answers to that. One, to some extent we do have to wait till the older of us die off. I'm sorry, you know history, and philosophy of science, anybody ...yeah, okay. So I'm not just being terrible.

Number two, the solutions that AI brings are definitely more, I think at this stage, at the granule level. So while there are governmental level, or state level conversations, that need to go on, and you're probably not going to have AI bots in there having those dialogues, but when it gets down to who's gonna implement the mandates, in all of these areas, again that's where you come in, and that's where, again, why the Digital Blue Helmet platform will allow tech folk to come in, either seconded by your organization, 'cause the UN's a place for organizations, rather than individuals, or individual volunteers working with a group of people.

Then when there's a specific project established, and how that gets established, and those things get chosen, and picked up, is still the gray area. But in my experience at the United Nations, there is a lot that you can do if you get in there and do it. And it won't happen if you don't.

Virginia Dignum: [inaudible 01:29:21]

Karen Judd S.: It's pretty basic, you know, it's not, as I say, it's not rocket science. And I had a third one, but it left.

Virginia Dignum: So, one more last question.

Audience Member: Just a comment, more than a question, I think. In response to what you've said. Just for what it's worth, I'm an entrepreneur, and I've been running a healthcare foundation that I started in India about three years ago, and the idea is how to use innovation to improve healthcare for the poorest people. We teamed up with the World Bank to identify the four or 500 most interesting healthcare innovations in India, and the issue was how do you actually get them out into the field, in remote areas, and urban slums, and so on.

You're exactly right, that a lot of the innovations that looked great at the developmental stage, at the prototype, and the lab stage, don't work that well when you actually get 'em out into remote, you know, high crime, low health indicator, type areas, very high poverty, and so on. But there is a process, which organizations like the UN, and the World Bank, and others have developed, step by step, to see what happens. In our case, I mean we are beyond just looking at innovations, we now run about 300 health clinics in remote parts of India, we're treating 10,000 patients a month.

So what that does is give us very large [desbeds 01:30:36] for using the most promising of these innovations, really seeing how they work. Demonstrating their efficacy. And once you demonstrate that, then governments get very interested in getting them in. So that's kind of the process, and I'm sure the UN has something similar.

Karen Judd S.: And there's part of the challenge, also the additional challenge is that when you do have something that works, on the ground, even if it is in remote areas, and I've worked with organizations that work on countering violent extremism, at which it requires a lot of education working with, you know, your key partners, your key peers in the religious communities in madrasas, and things like that.

Even when it works exceptionally well, at those levels, bringing it to the attention of the United Nations, the member states in the United Nations, and getting them to buy in and scale up that kind of work that is possible, there are always funding issues, there are always interest issues. It's a complicated world, it's not a simple one, but it is possible, and when those connections can be made, at key and central levels, then you're really pivoting some major stuff.

Anyway, but it really needs you guys. Just saying.

Virginia Dignum: Good.

Karen Judd S.: Just saying.

Virginia Dignum: Thank you very much. Also, thank you for all the speakers.