AAAI/CCC Symposium on AI for Social Good

Talk Sessions 2: AI for Social Welfare
Session Chair: Dr. Eric Rice

Fei Fang: Let's start our second technical session on social welfare and Eric Rice, who is also getting his ... Who also got his PHD from Stanford, and is a associate professor in the school of social work at USC will chair the session. So let's welcome Eric for the opening talk.

Eric Rice: I'm sorry that you all have to listen to me twice in one day. It seems a little bit unfortunate. But I want to pick up on some of the threads of the conversation that we've had already today in this ... In framing the idea of a social welfare technical session.

So one of the things that ... I actually threw in this slide about just two minutes before the coffee break because of what [Millen 00:00:58] said. So, brand new slide. Social work is very concerned with ethics, and I think that this is a really interesting issue about, you know, as AI is struggling with what should ethics be? This idea of art, I think is a really interesting one.

So these are some of the things which social work and social welfare research and practice is really concerned with. So service, you know that may or may not translate. Social justice probably does. Dignity and worth of people, the importance of human relationships I mean, there's a lot of social work practice that is about interacting with actual people. So again, this is maybe the importance of human computer interactions and relationships, perhaps.

Integrity and competence. So, you know, are you doing the best that you can, and are you doing it well? I mean this is sort of like in the Hippocratic Oath, the first "Do no harm" piece. This is the social work equivalent of that, is, "Try to do the best that you can for the people that you're working with."

I think sometimes something concrete is more useful than something abstract. So in practice, these are the major areas of research and practice that are focused on at the USC School of Social Work. Schools that have really good schools of social work and really good schools of engineering aren't always in the same place. Sometimes they are, sometimes they're not. Stanford, for example, doesn't even have a school of social work.

But these are the things we focus on at USC. So military social work, really issues around veterans and also around trauma and traumatic exposure that people have when they go on deployments to fight overseas. We're really interested in how to make social service organizations work better, so this could very easily translate into operations research sorts of problems, I think.

We're very interested in child welfare, so that really means both the system of intervening with children who've been neglected and abused, but also understanding what the consequences of abuse and neglect are, as well as the sequelae of abuse and neglect.

Then we're also interested in again. There's a bunch of people that are working on healthy aging. A lot of this is focused really primarily around healthcare access issues.

Homelessness is another area. So sometimes I think people think of the Center for Artificial Intelligence as the Center for Artificial Intelligence and Homelessness, because we've got so many homelessness projects going on. But that's not the case. We're diversifying our portfolio all the time, but we have a lot of people who do work on homelessness, because that's one of the foci of research in social work.

Serious mental illness and then behavioral health. So that's really preventable disease and substance abuse. So this really ties in nicely with the last session topics. Those public health topics are very much things which social work research is also interested in, particularly in so far as we're thinking of behavioral interventions for people.

So, another thing to think about a little bit when you're thinking about the intersection of social welfare work and computer science is, what is it that social welfare researchers and social work researchers, social work scientists, what do we do most of the time?

Most of the time, we do a lot of our work is survey-based research. So we collect large-scale surveys of high-risk populations to try to understand them. Like I said earlier today, we're really interested in interventions. So the idea is not just to do surveys of populations to understand them, but it's really to understand what are the factors that might lead to risk or resilience for particular issues?

So for example, what is it that might lead a kid to become homeless? Is a sort of problem that I've worked on in the past.

And so those can be community-based projects, administrative data, and sometimes there's publicly available data sets that we use. So I was really fascinated by one of the last talks where they were talking about cleaning and sanitizing data. Sorry about that. I guess Stanford wants me to register.

So, couple of other thoughts about social welfare research and then I will turn it over to our panelists. We oftentimes organize our thinking by populations. So, "Who is it that we're interested in serving?" Can be something that helps define the work.

And then other times, we define our work based on what is perceived to be a specific problem in the world, some sort of social problem. So suicide, healthcare access, violence and abuse,

But you can see that these are not sort of purely orthogonal dimensions, right? That I'm very interested in homeless youth as a population, but I'm also interested in the problem of homelessness. So where does that become a population, where does that become a problem?

And I think what I want to leave you as a parting thought before I let the speakers come up and speak, is that I've been asked at times what social problems we think are viable targets for this kind of research. And my answer is really any social problem.

The issue is more finding the point where there's something, where there is a problem within an actual social problem or within a research paradigm where AI can start to help create new solutions. This is what I think is the key, not so much that homelessness is very amenable to AI, but HIV or death and dying is not. I think in some respects, you know, any of these problems are equally amenable.

And I think one of the things that we have found to be very helpful, and I've heard this said several times today in different ways which is really nice to hear as one of the outsiders in this space, is that you all as a group at least seem to very much appreciate domain experts. And I would argue it's useful to have domain experts that are people who are boots on the ground, but also domain experts that are social scientists that have been working on those problems conceptually.

I think that one of the things that we've had with our program with homeless youth so much is that we have domain experts who are community collaborators that have been my partners for years and years and years, but then we also have a group of social scientists that are engaged in the work as well.

And I think that it's created something that seems like we might be a little less prone to some of the seemingly absurd examples that have come up, but obviously, you know, we also make mistakes all the time as well, too. So it's not like just working with a social scientist is going to solve your problems.

So, with that, let me begin to introduce the panel. So the first talk ... I'm not actually texting, I'm pulling up my, the program because Amulya told me that it had been changed. So first we are going to hear from Bryan Wilder is going to be presenting work, Uncharted But not Uninfluenced: Influence Maximization with an Uncertain Network. And while he gears up, I will take any comments or questions about what I just said, if you want to pipe ... Does anyone? No one?

Speaker 3: So what's hot right now in the social work research? Like which are the kind of hottest topics, I guess? Cuz in computer science everyone's talking about deep learning kind of thing, I don't know if there's any trend?

Eric Rice: Right, so the new trend in social work research that people talk about a lot is what we refer to as implementation science. So the idea is that we have a large number of interventions that we have done random clinical trials on over the past several decades. And we know that we have a good intervention for HIV prevention, or we have a good intervention for some aging population that's trying to get access to healthcare.

But the issue is really, "How do we bring them to scale?"

So how do we take them outside of a university setting, where they have -- there was a big National Institute of Health grant that funded it, and then put it into a community setting which is much more resource-poor and much larger.

And this is I think a great space where AI can really help because the scaling is difficult and cost is an issue. So if we can have models that sort of discount hypotheticals that we don't want to spend our time and energy on, it could be a great gain for us.

Bryan Wilder: Thanks.

Hi everyone. So my name is Bryan Wilder, and I'll be talking about some of our work on influence maximization and its application to HIV prevention.

So the domain is something that you've heard from Eric now about already. We're particularly interested in HIV prevention among homeless youth because of the enormous rates of HIV prevalence among that population. And so, as Eric talked about earlier, shelters will conduct interventions where they teach youth to be peer leaders and communicate with the others in their social network about HIV and encourage them to get tested for HIV regularly and things like that.

And the real motivation for this is that we don't have the resources to deliver interventions like this to everyone in the population. So we work through peer networks, and through social influence.

And this is where we have and opportunity as AI researchers, because there's an algorithmic problem hiding in here, that we're interested in which youth should we pick in order to maximize the dissemination of this information through the social network?

And this is something that's been studied under the banner of the influence maximization problem. So the idea is that you have a directed graph representing the social network, where the youth are nodes and the edges are friendships between them. And I have the ability to sort of at the start pick a set of nodes that are my seed nodes. So those are nodes that are active originally, have received my message.

And then in each time step after that, every node that's active has some probability of being able to convince their friends to become active as well. So that's how we model the spread of social influence through a population.

And we'd like to then select the set of initial seed nodes that, in expectation, maximizes the total number of activated nodes at the end of the process.

So this is a problem that's pretty well understood just as I've described it in the sort of classical setting. But when we think about moving these algorithms into more of a field deployment, there are new issues that come up that are, actually bring up what I think are really interesting new research challenges as well.

And so the first of these is that our domain is characterized by a lot of uncertainty. So there's a lot of missing observations, a lot of noise in our datasets. And what that means is that there's a lot of edges, right? That I might not be sure whether or not they exist. Or there's a lot of uncertain parameters about the problem. So classically in influence maximization, you're given a label on each edge. That's the probability that influence will spread. Like if you convince one friend, what's the probability that they'll convince the other? And of course, no dataset like that actually exists, where someone will give you those probabilities.

And then another feature of our domain, which maybe helps to make up for some of the uncertainty a little bit, is that it's really an adaptive problem. So as you conduct these interventions, you're able to learn more about the problem by asking the peer leader. So you can say, "Hey, this node, are you actually close with this other node?" And they'll tell, you know, yes or no or how often they interact, or something like that.

And so this is nice in that it gives us more information, but it also poses a new technical challenge, because you have to then reason about what's an adaptive policy that you should take instead of just a single shot decision.

So in this work, I'll introduce a new algorithm that we developed called DOSIM, which simultaneously handles both of these challenges: adaptivity and uncertainty. And we'll see that DOSIM can substantially reduce the vulnerability of our solutions, the different unknowns about the problem.

And then in the next talk, Amulya will be presenting some field tests where we deployed DOSIM and two other algorithms in the field, so stay tuned for that.

I'll start out then by introducing at high level how we modeled these two different problems. And so the way we think about adaptivity is that there's some prior distribution over the unknown parameters. So each edge has a prior distribution associated with the propagation probability on it.

And then, so there's some draw from that prior distribution that I don't get to see at the start. But then, as I interview nodes, they give us more information about their ties. So if a node says, "Yes, I'm actually like really close with that person." Then that shifts our posterior to a prior distribution maybe that has a higher propagation probability as sort of its average. And so that's how we model the sort of information that we gained over time.

And then the way uncertainty and robustness fits into this picture, is that of course I don't know what those prior distributions look like, so I don't know even on average what's the likelihood that influence will spread across a given edge. And so we can think about then as the prior distribution itself having unknown parameters. And then we have just uncertainty over those. So maybe I can get a very rough characterization that just tells me the average probability of influence spread is between such and such. And that gives us a set of possible prior distributions, any of which is equally consistent with what we know at the start.

And so our challenge then is to develop solutions that perform well, regardless of where the truth actually lies within that interval.

And so the way to think about this then is that we have a fairly expressive model that has these two sort of layers of uncertainties. So at the high level we have the sort of average behavior of the model in terms of what the prior distributions look like and that is itself unknown and subject to interval uncertainty. And then we have uncertainty dealing with what the realization of the draws are from those prior distributions, and that's something that we can sort of learn more about through qualitative interviews with participants as they become peer leaders.

Now, so previous work in this domain was the HEALER algorithm, developed by Amulya and collaborators. And so this deals with just one of those layers of uncertainties. So it assumes that I tell you what the propagation probabilities are on the edges, but some of those edges are uncertain and may or may not exist, and their existence is revealed as I choose nodes as peer leaders.

So this isn't thinking about parameter robustness, but it's handling the adaptivity portion of the domain. And the way it handles adaptivity is by modeling the problem as a POMDP, and which, since of course solving POMDPs is usually computationally expensive, it breaks the social network down to the sort of sub-communities, and then solves the simpler problem on each of those.

And then how DOSIM works in contrast to that, is that it handles the sort of basic level of adaptivity just via a very simple greedy policy. So it says, "At each time step just pick the nodes that will give you the best immediate return in influence spread."

And we see that experimentally this works very well, very comparably to the POMDP, but it's a lot more scalable. And that lets us use it as a building block in this bigger algorithm that addresses parameter robustness as well.

So now we'll transition over to how we think about robustness. And so the way we formalize the robust optimization problem is as a zero sum gained against nature. So I can think of it as that nature chooses a set of prior parameters i.e. nature chooses what's the average probability that influence spreads across all of the edges.

And then the algorithm has a set of strategies which consists of policies for choosing seed nodes. So in response to any set of observations, what's the set of seed nodes that I'll choose? And then the payoffs to this game are the ratio comparing what, in expectation, will my algorithm get versus what would the optimum be if I were told nature's draw? So if I knew what the true parameters were, how well could I do? And we'd like algorithms that perform well so this ratio is pretty high, regardless of what nature decides to do.

And the immediate problem then, and the reason that we can't just solve this game is that the strategy spaces are enormously large. So nature has a continuous strategy space where you can pick any probability within a given interval, so that's not even a finite strategy space. And then even the algorithm has an exponentially large strategy space, all the different policies for picking seed nodes.

And the way we handle this is to take a double oracle approach that incrementally builds up the game. So instead of trying to solve this attractively large game to start with, we start with a much smaller game and repeatedly add strategies to arrive at an equilibrium.

So here's what this might look like. I start with a very small number of strategies for each player, so just a couple of policies for the influencer, and a couple parameter settings for nature. And then I can write down the game matrix here with the payoffs associated, just with those strategies. And then it's very easy to write, I can write an LP to find the equilibrium of this game, because it has only a handful of strategies.

Then, I have an oracle for each player, which is a best response oracle. So for each player, this will tell me in response to the strategy just on this very little restricted game, what's the strategy that this player would take and best response. So for nature this would say, "Given these two policies, what's the set of parameters that sort out the toughest case for the algorithm right now?" And then will add that set of parameters to the table.

And then the influencer has a similar oracle, which based on greedy policy says, "What policy would you take in response to these different sets of parameters?" And you add that to the table. And you keep going like this until you get to an equilibrium. And so we see experimentally that the algorithm finds equilibrium solutions with very sparse support, so it terminates just after maybe ten or 15 iterations in practice.

So I'll give a very brief highlight of the theoretical analysis of this algorithm. The essential question that we want to answer in this analysis is that since influence maximization is itself a very hard computational problem, is it sufficient for me to be able to give you a sort of approximate best responses for the influencer? So if I can you an approximately good policy for selecting seed nodes, given fixed parameters, does that translate into approximately robust solutions?

And so we answer this question affirmatively. We show that if you give us an alpha approximation algorithm for influence maximization in whatever your domain is. If you give us an alpha approximation algorithm for that domain, then our algorithm will in turn give you arbitrarily close to an alpha approximation to the optimal robust policy.

And so this is a nice property, because it means that since we can easily verify that our algorithms are effective at performing influence maximization, just looking at the datasets comparing them to the optimum one. It can be characterized. Then this tells us that we're actually close to the optimally robust policy, which is something that's much harder to sort of just tell by inspection.

And then I'll also give just, again, a brief idea of experimental results. So this in simulation, Amulya will present the sort of real results later. But this is in simulation on networks that were collected from homeless youth at two drop-in centers in Los Angeles. And we compare two algorithms, so DOSIM and a greedy algorithm that just optimizes based on a fixed set of parameters that we sort of had from previous work as sort of the best guess.

And then we say, "What's, in the worst case, what percent of the optimal influence spread will these algorithms get?"

And so we get that DOSIM is always within around 90% of the optimal value, where if you just plan based on a fixed set of parameters, like people ... Previous work had usually done. Then you can lose, you know, a fairly substantial amount. You can lose up to a little bit over 30% of the optimal influence due to parameter uncertainty.

So to wrap things up then, what we found was that when we tried to move into more of a real world field deployment setting for influence maximization, there are substantial new technical challenges, really core research challenges. And that in particular handling uncertainty in a thoughtful way can be very important. And so with that, then I'll turn things over to Amulya who has some results about field tests.

Eric Rice: So do we have a question or prior to, while Amulya sets up his presentation?

Speaker 5: Excuse me, I have a question. Could you give us some intuition about your result? Is there something like some modularity or?

Bryan Wilder: Right, right. So modularity is, when you're thinking of sort of the lower level of uncertainty, so just what should I do for fixed parameters? So our problem interestingly isn't quite so modular, because the information gain essentially sort of messes with that structure. So the result that I presented theoretically is sort of at a higher level where I say, assuming that you can solve whatever the underlying combinatorial problem is to some level of accuracy, then sort of in the outer loop, "What robustness do we have?"

And that's based just sort of analyzing any ... And that applies to sort of any generic game, right? Where you can give me approximate best responses for either player. So the reasoning for the theorem that I showed you is sort of at a higher level of abstraction than the domain itself.

Speaker 5: And then relating to it, may I?

Eric Rice: Sure, sure.

Speaker 5: How large are your networks then? Because, for example, we actually looked at influence maximization for a different setting, you know? And we realized for example, even for [inaudible 00:22:07] problems where you use greedy, but greedy would actually take forever to run because you really, at each node you would have to decide it. So do you have a similar situation here where you have ... I mean number one, I don't know how large the network is and the greedy might still be quite time-consuming.

Bryan Wilder: Yeah, yeah that's a good question. So for our domain the networks are relatively small, you know maybe 150 to 300 nodes or so, so it's not really a problem for us. So I'd say we've thought about this problem a little bit, you know what would we do if we had to handle larger networks? And so there ... I think like from an algorithm design standpoint it's hard to hope to get that much faster than greedy.

You know so like maybe where you think about speeding it up is that the computational expense lies in evaluating the influence spread, and there there's sort of a lot of clever ... Like you know, there's a big literature on sort of clever tricks too, if you can't sort of afford to do that in maybe the most expensive sampling fashion. Then how could you speed up that process a lot?

And so generally I think the really scalable algorithms for influence maximization that can handle like, you know, several hundred million node graphs. They retain the same greedy approach but they speed up the evaluation of the objective.

Eric Rice: You had a question? Cuz it looks like Amulya's still got a minute to get his thing going.

Speaker 6: Yeah, I was just curious about your modeling approach. So you have this robustness framework because you don't know what to make the priors for the link probabilities be. Like maybe the more conventional Bayesian approach is to do some kind of hierarchical model or something like that? Where you put a prior on priors is ... Like what motivated your choice, your modeling choice there? Was it computation or do you think the robustness is really essential? Or both?

Bryan Wilder: Yeah, so it's certainly not computational because having a hierarchical prior would be very cheap computationally.

Speaker 6: Okay.

Bryan Wilder: But, yeah. So from our point of view, I mean when you say put a prior on the prior, I mean if we don't know what the prior parameters should be, we certainly like, we don't know what the hyper parameter should be. And so maybe the hope is that they'll sort of have less influence as you go up the [crosstalk 00:24:14]

Speaker 6: Right, that's the idea, right?

Bryan Wilder: But it seems better instead of just introducing higher order problems about how to quantify uncertainty, to see if we take a more conservative approach, right? And say that it will work just for anything in this interval, then can we still find good solutions? And that's sort of a stronger guarantee I think.

Speaker 6: Okay, thanks. I was just curious cuz it's kind of a little outside the usual Bayesian Paradigm, so thank you.

Bryan Wilder: Yep.

Eric Rice: Amulya are you just about ready?

Amulya Yadav: Yeah.

Eric Rice: Okay, so let me introduce Amulya Yadav, who's going to be talking about Influence Maximization in the Field: The Arduous Journey from Emerging to Deployed Application, which I've been a part of.

Amulya Yadav: All right. Hello everyone. Today I'll be talking about deploying influence maximization algorithms in the field. This is joint work between our group at USC and Eric's group at the USC School of Social Work.

So following on from Bryan's talk, so far we now know that homeless youth are extremely prone to HIV, and there are these homeless shelters who conduct these intervention programs to raise awareness about HIV amongst homeless youth.

Now the shelters want to strategically pick key influential homeless youth in these social networks and train them as peer leaders so that these peer leaders can then raise awareness about HIV amongst their peers in these social networks.

Now a key computational question that these shelters face is, "How do you select these key influential homeless youth? How do you select your peer leaders?"

And in previous work, HEALER and DOSIM were two AI-based algorithms that were developed by a group, and in simulation they've shown good results. And now we want to deploy HEALER and DOSIM in the real world.

So the first question that we want to answer is, "Do HEALER and DOSIM perform equally well in the real world?" And we aim to answer that question by conducting pilot studies to test HEALER and DOSIM's performance in the real world.

We have two goals in mind when we do the pilot studies. First we want to verify that these AI algorithms are indeed needed in the real world, that it is not the case that simple heuristics like degree centrality or some other centrality-based measures can give you just as much.

And secondly we want to ensure that these algorithms are indeed usable in the real world, which is very important before these algorithms can be deployed on a large scale.

These are the first studies which compare algorithms for influence maximization in the real world. So these are the contributions. We conduct three pilot studies with 173 homeless youth in Los Angeles, which provides us with an opportunity, for the first time, to do a head to head comparison of different influence maximization algorithms.

We analyze the results from these pilot studies to understand why is it that simple algorithms fail to perform well in the real world, whereas the AI-based algorithms, HEALER and DOSIM, perform well. As a result of all this, we are able to raise awareness about HIV amongst homeless youth using AI-based methods, which is a first.

So for all three pilot studies, we go to a homeless shelter. We work with two different homeless shelters: Safe Place For Youth and My Friend's Place. We go visit these homeless shelters and for each pilot study we've accrued approximately 60 homeless youth. And then we are gonna select, out of the 60 homeless youth, we are gonna select four peer leaders for three successive interventions.

And for the first pilot study we are going to use HEALER to select the peer leaders. For the second pilot study we'll use DOSIM, and for the third pilot study we'll use degree centrality to select these peer leaders. And then we'll compare what happens.

As an example, for the first pilot study, we went to Safe Place for Youth. This is in Venice Beach in Los Angeles. And we've accrued approximately 60 homeless youth from the shelter. And this is how the network of homeless youth looked when ... The network of the 60 homeless youth looked like. HEALER looked at this network and came up with a recommendation of these four nodes to select as peer leaders in the first intervention. And then in the second intervention it came up with these other four nodes, and then similar things happened for the third intervention as well.

Now in this network there are these black nodes which are the peer leaders whom we have directly influenced. And there are the rest of the nodes, which are the non peer leaders, whom we have not directly influenced.

Now in these pilot studies, what you want to measure what fraction of these non peer leaders actually get informed about HIV by the end of these interventions. And that is what we'll measure, I think you've seen this [inaudible 00:28:49] slide as well.

So the Y-axis is showing the percentage of non peer leaders that are informed about HIV by the end of these interventions, and as you can see the peer leaders selected by HEALER and DOSIM were able to spread information to approximately 70% of the non peer leaders in the network, whereas degree centrality was only able to spread information to 27% of the non peer leaders.

Now the people in the blue, they're the people who got information about HIV. What percentage of these people actually started adopting safer behaviors? And so, you know, was there a behavior change? And that is what we'll be measuring next.

So of these people in the blue who got the message, in HEALER and DOSIM's case, approximately 30% started getting tests for HIV by the end of this interventions, whereas surprisingly in degree centrality, none of the leaders, none of the non peer leaders who were informed about HIV had started getting tests for HIV regularly.

Now why is this happening? Why is degree centrality performing so poorly despite it being such a natural principle to be used? When we analyzed the real world networks that we were playing with, the real world networks that were being used in the pilot studies, we realized that, you know, most of them had a lot of community structure. They were composed of tightly knit communities with very few edges that were going in between the communities.

So for example, this figure shows the percentage of edges that are going in between the communities across all three pilot study networks. As you can see, across all three networks, approximately 13 to 14% edges go across these communities, so they're fairly disconnected.

Now since within a community there are a lot of edges, so they are very densely connected within a community. What that means is nodes within a community will end up having similar degrees. So, you know, they will have similar number of connections. This means that degree centrality-based approaches would focus their efforts on just a single community or just a couple of communities, completely ignoring other communities in the network.

Indeed this is what we saw in the pilot studies that we conducted. Degree centrality, in all three interventions, was focusing its efforts, was picking all its nodes from just a single community or a couple of communities. Whereas HEALER and DOSIM were able to spread their efforts, were able to diversify their efforts, across different communities.

So this is one reason why DC's performing poorly. The second reason is that in a network, when you pick peer leaders, there are going to be many edges that go in between these peer leaders. Now these edges are redundant from an influence maximization setting, because information spread along these edges does not matter, because both endpoints are already influenced.

And this figure is showing that, from the picks that degree centrality made, it ended up creating approximately 27% redundant edges in the network. So 27% of the edges were no longer used for influence maximization, whereas HEALER and DOSIM, the picks that they made, they ended up creating less than half this number.

So to summarize, degree centrality is performing poorly because A: it fails to exploit the community structure, and B: it generates lots of redundant edges. There are many more reasons that we've outlined in the paper that we wrote, and please feel free to go read that.

So to summarize, this was the first study which compares influence maximization algorithms in the field. We conducted pilot studies with 173 homeless youth in Los Angeles. We found that the AI-based methods are actually providing value. They are outperforming simple degree centrality-based methods significantly. As a result, enthused by this result, we have started, actually, a much larger study with 900 homeless youth to test out many different, more sophisticated baselines against which we'll compare our algorithms.

That's it. So I want to end by showing you a video which shows the impact that these algorithms have had on the lives of these homeless youth.

Let's see if I can get ...

Eric Rice: (on the video) The only more exciting than the quantitative data that we collected is the unexpected human impact that this program has had on the young people who are the actual peer leaders in the intervention.

Michelle(video): I think specifically with this population, where a lot of the experiences are kind of being invisible or being not acknowledged to give them an opportunity where their voices can be heard.

Cody (video): Naturally I'm an introvert, but being able to be a part of this, this group. It allows you to actually reach out to people, you know. To make it feel like you actually have knowledge that you can give to someone.

Alison (video): I think [inaudible 00:33:27] now has a language around what he wants to do. He has a language about about affecting change, and being a leader, which I don't even think he knew what that language was.

Blue (video): I know I'm a goofy person, and I know I have talents, but I'm very goofy. And it's hard for me to control my goofiness when it's time to be serious. So when they told me, I was like, "All right."

Eric Rice: (on the video) When we pick these young people and give them the opportunity to be trained to be a leader, we could see that it was really changing the self esteem that these young people had, and the sense of confidence that they could be an agent for positive change in the world.

Blue (video): A few of my friends told me that they look at me as a leader, so that was ... It felt good, but it was kind of weird.

Amulya Yadav: That's it. Thank you.

Eric Rice: Any questions? Yeah. So we have, uh ...

Speaker 12: I want your producer.

Eric Rice: Huh?

Speaker 12: I want your producer.

Eric Rice: This is one of the benefits of being in Los Angeles, you can find people who can do things like this pretty nice. [crosstalk 00:34:32]

Speaker 12: So to come back to the methodology, I have a question about this. So I feel like we've known degree centrality is really horrible for representation in social networks for like, I don't know, 40 years. So like, there's other methods, right?

Amulya Yadav: Definitely.

Speaker 12: Like betweenness centrality or eigencentrality and all of the stuff that Kathleen Parley's group has been doing for years. So why wouldn't you choose one of those as a baseline to compare to?

Amulya Yadav: So I think the knowledge, the sort of understanding that degree centrality is poor, that is, I believe, more common in the computer science community than it is in, you know, the social sciences. I mean people are still ... I mean when you talk to the homeless shelter officials, for example, they, despite whatever reasons we have given them, they were not convinced why wouldn't picking the most popular people make sense?

And if we really want to implement this program with them, I mean getting past this barrier that, "Look, the method that you have been using for all these years does not really make sense and we can prove it to you."

You're very right, we need to test out our algorithms with more sophisticated baselines. There are many more baselines. And so, you know. But in order to ensure that we are able to actually test these algorithms, we needed to get past this barrier first. So that is why this baseline was chosen as opposed to other baselines.

Right now, as I said, we are trying out ... We started a much larger study, and there we intend to use more sophisticated baselines to compare against. Yeah.

Eric Rice: [inaudible 00:36:06]

Speaker 13: Yeah, so I had the same, one question I had was the same, basically. As a baseline, and at least in the larger study, it seems to make sense to use something like eigen and/or betweenness centrality. But the other thing, so that's one issue with degree centrality.

The other issue that you've explicitly mentioned is I presume you just rank by degree centrality, right, and pick the highest ranked?

Amulya Yadav: Yes.

Speaker 13: Right. But from a computer science perspective that's also the wrong thing to do purely algorithmically. What you want to pick is key people with the highest joint degree, right? This is why [crosstalk 00:36:40] This is why you don't use redundancy. This also would be ...

Amulya Yadav: That's true.

Speaker 13: And this transfers to the other centrality measures or whatever you use. It seems like a baseline would be to use [crosstalk 00:36:50]

Amulya Yadav: So you'd pick the top kid and then you'd remove them, and then you'd pick the next top kid?

Speaker 13: No, no. You take the [crosstalk 00:36:57]

Amulya Yadav: I believe you'd get the same result, right? I mean if you take the top kid and then you remove them from the network. [crosstalk 00:37:02]

Speaker 13: No no. [crosstalk 00:37:02]

Amulya Yadav: And their edges.

Speaker 13: You don't ... You take the top remove ... Yeah, yeah.

Amulya Yadav: Yeah.

Speaker 13: I mean, you get the right idea, I think. I think it's just a communication may be confusing. But you basically take a key with the highest, let's say ... [crosstalk 00:37:16] Right, that's sort of straightforward.

If you're talking about other centrality measures there's sort of similar things you would want to do.

Amulya Yadav: Yeah. I mean I agree, yeah.

Speaker 14: I have a ... As someone that ... As a social worker that's actually in the field, you know, implementing these things. If I have no ... If I'm gonna implement one of these interventions, with no like access to any sort of computational methods, degree centrality is kind of what makes sense. It's easy to identify who's most popular just visually or by being in the space. So that's kind of why degree centrality makes sense for a baseline.

Does that make?

Speaker 13: Yeah, it makes perfect sense so this is, the fact that we're having this discussion is already incredible, right?

Speaker 14: Yeah.

Speaker 13: I mean [inaudible 00:38:02]

Eric Rice: Well and in the field of social work and public health, they've only ever done degree centrality, and typically it's through very qualitative methods. They don't even record degree and then rank order people. It's usually sort of impressionistic.

Like, "Oh, I've been watching the fact that Carla seems to talk to a lot ... You know, talks to a lot of people so she must be really popular, we should pick her." You know, sort of thing.

And there are people who have been suggesting that betweenness centrality would be, or eigenvector centrality would be a superior method to do this, but there's almost no one who's implemented that in actual practice.

Speaker 14: Yeah.

Speaker 3: In sociology or social work?

Eric Rice: In public health intervention contexts. So within sociology when you're doing just conceptual work, which is what Kathleen Carley primarily does, she can talk, and does talk about how worthless degree centrality would be. I mean people ... Actually here at Stanford, Karen Cook was talking about how degree centrality doesn't equal sort of power in situations. I mean this has been going on since the 70s. But it's hasn't translated into these more applied spaces.

And that's actually fairly because of the history of the University of Chicago kicking the women who started social work in Chicago out of their department, but that's a whole other story for a different day.

Speaker 5: I was just going to [inaudible 00:39:18]

Eric Rice: Sure, sure.

Speaker 5: I was just going to make a comment, because I think that's actually very important. Not only do you have to do good work, you also need to gain the trust of the people who are implementing your algorithms and, you know. And I think that is actually really really important.

I mean my own experience in a completely different context, but I actually did, I worked for an airline doing [inaudible 00:39:48] scheduling. And I remember when we would produce schedules, you know, they would not really trust us. So we actually had to get the people involved, starting with their own algorithms. And so I think this is actually very nice.

And I think you should write about that. About, you know, this little nugget that you really need to have solutions that will gain trust, and how to go about that from this social work. Cuz I think that's actually a very interesting point.

Amulya Yadav: Thank you, thank you.

Eric Rice: I need to pull out my phone again but not cuz I'm gonna text. Thank you. Now next up we have Ben Ford, and let me just pull up the title of your talk while you're pulling it up, which is going to be Cloudy With a Chance of Poaching: Adversary Behavior Modeling and Forecasting With Real World Poaching Data.

Ben Ford: Hi everyone. Thank you for attending, and yeah. So I'm going to be talking about predicting and preventing wildlife poaching with real world data.

So wildlife poaching is a worldwide problem among the many well-known species such as tigers and elephants. They're being poached in large numbers to meet the demands of the wildlife trade market. The illegal wildlife trade market.

And our ultimate goal is to use basically our predictive analytics, be able to predict where poaching is occurring in order to better inform ranger patrols to basically, so they can go out and basically be more efficient with their limited resources and ultimately prevent wildlife poaching from occurring.

So the first problem is how do we actually accurately predict where poaching is going to occur in the real world? Real world data poses a lot of unique challenges. In addition, once we predict it, how do we actually evaluate it in the real world? We all know simulations and real world evaluations are very, very different things.

So we're working with rangers in Queen Elizabeth National Park in Uganda. It's a park about 2000 square kilometers, and they've given us data from 2003 to 2015. And this dataset consists of various geo-spacial features such as is an area foresty, does it have marshlands?

Also various distance features as well. How far away is this from the nearest ranger outpost, for instance?

Also there's also how often has this been patrolled? So you can see Fei and Melind are actually in this picture in Southeast Asia.

And also has crime previously been observed here? Have they found snares, poached elephants, etc.?

And some of the key contributions of this work is that we did pretty a pretty extensive evaluation on this 13 years of real world data, and we found that really surprisingly to us, is that a decision tree ensemble, relatively simple model, outperformed other like previous state of the art models that were really complex, which I'll describe capture in the next slide.

And in addition, since it's a decision tree ensemble, it can be trained very quickly, which is good for rangers with limited computing power. In addition, it can give us interpretable rules that also helps to build trust with the rangers and domain experts to really validate what our model is doing.

And as a result, they actually agreed to have us deploy our predictions in the field, and they were actually able to find, in a single month, about ten times the number of findings that they usually find in a given month. They found a lot of snares, they found a poached elephant, and they also found a lot of of signs of trespassing. And basically in collaboration with both the Wildlife Conservation Society and Uganda Wildlife Authority, might have actually helped them to save some animals from being poached.

So the previous state of the art capture is a two layer model that basically takes care of the time dependency between poachers actions. Basically what they do now informs what they'll do later.

Also, the fact that rangers can't observe crime perfectly. If they patrol an area, they might have missed that snare because it was in dense grass. They might have just walked over it and incorrectly labeled that area as being not attacked.

This model attempts to model all of that with a logit function, and thereby also it gives a ranger observation probability. It's like, okay, so given these factors, given these input factors, basically those features I described earlier, what's the probability that a ranger will observe an attack at a given location?

However, since it's a logit model but using a lot of different features, it wasn't very interpretable, whatever output it was giving for the model. Additionally, it took a very long time to train, which is not acceptable with limited computing power in Uganda.

So I'll just briefly go over some of the key features of our decision tree ensemble, then I'll discuss our empirical evaluation and real world deployment results.

So our decision tree ensemble is a standard ensemble. We have a majority voting mechanism. Each tree is trained on about nine different features. Basically that's all we really have access to. Also each tree that got trained has an average depth of about nine, so we tried to make sure that it wasn't too deep.

Also, we found that a five tree ensemble will perform best. It kind of gave the best trade off between still not being too too large, like we don't want to have a hundred tree ensemble. But also it performed better than like, let's say a three tree ensemble.

So basically we trained ... For the purpose of this presentation due to time I'm only gonna show results for one particular training test combination. But we're basically trained on 2003-2014 to predict where wildlife poaching would occur in 2015. And it presents some pretty standard metrics. Precision, recall, F1 and L&L, which I'll go over in the next slide as well.

And overall, we tested about 192 different model configurations. I'm obviously not going to show all of those now, more will be in the paper.

So here each bar corresponds to a different baseline that's been tested. So blue corresponds to if we predict everywhere, positive baseline. Uniform random, 50-50 chance, capture, superior state of the art. And the last, the purple bar is our ensemble. As you can see, even though the ensemble doesn't have as good recall as the other methods, its precision is a lot higher. So for this last metric, L&L, basically what happens is that if it has really good recall but a lot fewer predictions to get that good recall, it will be a reward more heavily than something that just predicts everywhere and gets good recall.

So for the real deployment, basically asked rangers to patrol two 3x3 square kilometer areas that were infrequently patrolled. We didn't want them to patrol somewhere where they knew what was going to happen. This is also encouraged exploration of areas as well. And we predicted it to be attacked. And as you can see, they found a lot of signs of trespassing, lot of different snares as well too, in addition to illegal fishing and plant harvesting.

Trespassing, here you can see campfire ashes, basically signs of litter. Indirect signs. They also found a poached elephant, and also snares as well, too. Here's a picture of a snare that they found while on patrol.

And we have some ongoing experiments. Basically this is more of a park wide experiment. You can't really see the colors too too well, but basically we divided up different patrol areas into three groups, where we have a group one, which is basically, "We predict this area's going to be attacked a lot."

Green would be basically, "We're not predicting many attacks at all in this area."

And yellow being, "There's going to be a moderate amount of attacks happening in this 3x3 area."

And here are some pictures of the rangers actually training to use the data input devices as well and follow these patrols.

And we actually have one pretty big finding at this location. Rangers actually followed a trail that they found, and were able to actually ambush a camp and arrest one of the seven poachers that they found, which is a pretty rare occurrence. Basically I think in the dataset I found it was like less than 1% of all the observations led to an arrest. So that's a pretty big, pretty exciting finding. They confiscated a bunch of harvesting tools and wire snares and everything.

And in addition, they've also been finding indirect poaching signs. They've chased poachers out of the park. There were also signs of road building, which was kind of odd to us. But if you think about it, it can help poachers get in and out of the park easier.

But yeah, so that's basically the end of the talk. Yeah. I'll take any questions now, thank you.

Speaker 16: Can I go first? So you mentioned there's a fairly small set of features, so what kinds of features then turned out to be useful for these decision trees?

Ben Ford: Yeah, so we tried to do some principle component analysis to kind of see if we could like trim down the feature space a bit. And it turned out that most of the features ended up being very important. In terms of like explaining what's going on, that's also something that's ongoing research where we're trying to see, "Can we use a couple features to explain why the predictions make sense?"

Eric Rice: He's asking what exactly [inaudible 00:49:22]

Ben Ford: Yeah, so I'd say like, we also have ... There's various distance features. So basically how far away is this from like a fishing village for instance? So if there's going to be illegal fishing, you'd think that that would be important, this distance to the park boundary that we have now. How many animals are in this area?

And also other things too such as the terrain information, like whether or not it's heavy forest, in which case a lot of people don't go in there cuz it's just too difficult to travel in there. Or if it's just flat grasslands which are really popular poaching areas.

Speaker 17: I had a question about how your models respond to time.

Ben Ford: Yeah.

Speaker 17: So you're training on a ten year spans of time where poaching activity has evolved reacted to the ongoing efforts of the, you know, anti-poaching officers there. How does your model attempt to dilate between years or times of year or actually responses that have changed because they just changed how they were doing the anti-poaching.

Ben Ford: Yeah, that's a great question. So this current model kind of like does a, basically compacts the time and everything like that. So what we're able to predict is basically how attractive is a particular area? Which might not change as much over time, but we were also looking into basically training on different batches of time too to see if like, "What happened in the last three years? How does that impact poaching in this current year and everything?"

Eric Rice: I think we'll have to save any more questions for after just so that we make sure that everyone gets an opportunity. So our next speaker is Ayan Mukhopadhyay, and he's going to be talking about Optimal Allocation of Police Patrol Resources Using a Continuous Time-Crime Model.

Oh, here, why don't you have the mic?

Ayan Mukhopadhy: Thank you. So quick question, how are we doing the time? Are the cards you're showing, are they to the eight minute limit or the ten minute limit?

Eric Rice: They're to the eight minute limit.

Ayan Mukhopadhy: Okay, great.

So good afternoon, everyone. I am Ayan Mukhopadhyay. The talk that I'm going to present is titled Optimal Allocation of Police Patrol Resources Using a Continuous Time-Crime Model. I am from Vanderbilt, in the Computation and Economics Research Lab led by Eugene Vorobeychik. This work was done in collaboration with Professor Melind Tambe's group at USC.

And we've heard a lot about collaborating with domain experts. Professor Kenneth Pence is a social scientist at Vanderbilt. Sorry, Professor Paul Speer is a social scientist at Vanderbilt. Professor Kenneth Pence is now a faculty member at Vanderbilt, but he was with the Nashville Police Department for over 20 years so it brings a lot of domain expertise into how we look at crimes.

So we saw a lot of technical details in the poster session, at least some of did. Here we look at a very high level view of what we are trying to do with the project and how we are trying to do it.

So crime prediction and predictive policing are something, two things that are very closely related. So from the FBI Annual Statistics in 2014, those numbers give you an idea about why crime prediction is important. The total number of crimes, and more importantly the total number of potential targets, has a gross mismatch with the total number of law enforcement officials, and this is just in the United States.

So the end NIJ defines this idea of predictive policing as the future of law enforcement, where slowly all police departments are looking at at least some kind of a primitive crime mapping in order to understand where crimes would happen.

So predictive policing has two major consequences. One is proactive policing and the other is passive policing. So proactive policing looks at deterrents and then preventing crime, while passive policing looks at responding to calls, to crimes as fast as possible. And in this talk we look at passive policing and also how we develop models that help us achieve that.

So crime prediction is often confused with crime mapping. Crime mapping has actually been there for a long time, especially to plot locations of crime on a map and identify hotspots that have high frequency of crimes. So the goal of this project is to go beyond that. We want to create formula models that can forecast crime incidents in time and space, and then we want to use that model for optimizing police placements.

So as far as crime prediction is concerned, here is what we want to do. We want to find distribution of crime incidents in time and space. We want to allow for the inclusion of any arbitrary covarients. For example, we want to answer questions like, "If it snows tomorrow, what are the most probable locations where crimes would take place?"

Or, "If there's a new liquor store that opens in this area, how does crime shift?" And questions like that.

We want to estimate and learn our model from data in a principle way, and we also want to capture this notion of deterrents and the effect of police presence on crime. So once you allocate police, does crime shift? And if it shifts, how does it shift?

So the basic approach as far as looking at this problem spatially is concerned, is that we discretize space into equal-sized grids, and then we can learn one model for each of the grids or we can learn one model for all of the grids. In practice we do something called hierarchical clustering where we would individually merge grids, similar grids. But we'll not go into that. Imagine that there is a way to deal with it.

In order to look at the problem temporally, what we use is something called survival analysis. So for each grid allocation we want to predict time til the next incident. So in case you're not familiar with what survival analysis is, think of it as a parabolistic way of learning and estimating time to incidents. In its original form it was something slightly different, it was about estimating risk or hazard.

But there parametric models, especially the one that we use called accelerated failure model, where you can actually model time to incidents. So it would take time to incidents and a bunch of arbitrary covarients as its input and it would give you a distribution over time to incidents.

So f of t given a function of covarients is what we are trying to learn, and what this does is it provides us with a very natural way of including any covarient that we want, and then understanding its effect on it affects crime. And we can estimate the model using maximum likelihood estimate.

So the set of features that we use, we use a bunch of features from Risk-Terrain Modeling, and then temporal as well as spatial-temporal features. I'll quickly go into three of them. We use police presence, so how much police was present just before a crime. We use recent crime occurrences, so how does recent crime occurrences in an area affect future crime occurrences? And then crimes spill over from neighboring grids. So if police goes to a grid that is further away from where I am, criminals could shift and potentially come to different grids to commit crime.

So here are our results. We calculated, we compared our results with two state of the art methods. One is a dynamic base network method and the other is DSDA, which is a combination of time series analysis and a software called CrimeStar, which is now part of how NIJ looks at crimes.

And we found that our method was at least as good as those methods, and more importantly it was dramatically faster which helps us deploy real dudes.

So we can predict crime, now what do we do with that? We want to optimally allocate police resources to respond to crimes as fast as possible. And this problem has not really been explored much in literature. What we want to do is we want to minimize the expected response times to crime. So this is a dynamic optimization problem under uncertainty.

And what's interesting and what makes this problem tricky is once you're done with all this, you allocate police and crime shifts. So you have to be aware, come up with a way to address that.

So what we do is we treat this as a two-stage optimization problem. Think of it in this way that in first stage, police allocates resources, uncertain about where crimes would happen. Given that allocation crime happens, and then in the second stage police actually responds to those incidents. And the goal is to minimize expected response times.

So the way we do this is that we map this problem to two classical optimization methods. I'll go over the simpler one, the transportation problem. Imagine police as suppliers and crimes as consumers, and the service or the good you are transporting is actually police themselves, so they have to be, they have to move to crimes in order to respond to them.

So as I said, so we use Bender's Decomposition to iteratively add constraints and solve this problem, and we calculate the expectation by using Sample Average Approximation. So the challenge is, as I said, once you allocate police, crime moves.

So the way we deal with that is we came up with this iterative approach. So we intuitively give police repeated chances to respond to crimes. So you allocate police, crime shifts, you would reallocate police and you would keep doing this til it converges.

So we've done simulations and compared our approach with the actual approach that the Nashville police department has, and we found that with respect to both actual crimes and simulated crimes, we did significantly better than the actual approach that of dispatching police patrol vehicles.

So as part of our ongoing work, we are trying to create a unified model of prediction and response so that multiple types of emergency responders can work together, which will hopefully be decision theoretic approach. And we are also trying to create an app that multiple responders can carry. For example, police department and the fire department, so that they can communicate with each other, communicate with their home stations. They can visualize incidents, visualize where other responders are, and hopefully will work collaboratively to respond to incidents better.

Thank you. I'll take questions.

Eric Rice: Questions?

Speaker 17: One could imagine situations where you're just chasing after the crime if there are too many criminals and not enough police. You just keep displacing them and then you keep chasing after them so that it never converges. Cuz, is that possible if it's under-resourced?

Ayan Mukhopadhy: So, as I said, we formulate the inner optimization problems ... So as a transportation problem, which must be balanced, so you must have enough police to respond to crime. But if that's not the case, so the approach that I did not go about in the presentation but was in the poster, is that you could map this problem to a key server problem, where you have requests coming in and you have servers that need to move to attend to the requests, and you can [inaudible 01:00:33].

So it is slightly intricate because you would need an ordering of incidents. But since our prediction mechanism is spatial-temporal you can actually create an ordering. So you would know how exactly in what order incidents happen, and you could respond to them one after the other, if you don't have enough police.

But another quick point. We never observed that in our data set in Nashville. So we break our day into slots of four hours each, and so crimes are relatively rare events, at least in Nashville. And we always had more police than crimes.

Eric Rice: That's not like Los Angeles. Is there another?

Speaker 14: So again, as a social worker in the room and the non computer scientist. You know, predictive policing always kind of, it just makes me feel a little icky and like really nervous cuz it just feels like one of these technologies that might, if given into the wrong set of hands, it could be exploited or have a lot of bias, and, I don't know. How do you go forward with this kind of technologies and keep those things in mind? That there can be rates of error in the implications, and things like that.

Ayan Mukhopadhy: So as far as bias is concerned, I think if ... The first point. What we're looking at is date about reported crimes. So if there is a pattern in the data that has an inherent bias, I don't see how we can be agnostic to it. So it really depends on the data we have. If there is something which is an obvious part of the data, we will learn it.

But then ... So none of these methods I think are designed to absolutely replace the role that police has about, again, domain expertise. I mean we do not tell police who to arrest and who not to arrest, for example.

Speaker 14: Yeah.

Speaker 13: So actually the choice of the objective I think is important here as well, for the reason that you mentioned. Incident response, the fact that we're doing prediction as a means to do incident response, it seems to alleviate a lot of these concerns. We're not trying to, you know, target anybody or arrest anybody. We're just trying to get to people who, whenever something bad is happening as fast as possible so we can help them.

But close world assumption is of a course a problem here as well.

Speaker 14: Just something like that, if it was even framed and said, "incident response times." Like, "increasing crime response efficiency," like makes me feel a little bit better than just the word predictive policing or things like that. Crime prediction, cuz that kind of weighs a little bit heavier. Even just in the way.

Ayan Mukhopadhy: So also, I think at some point this helps us eliminate some of the bias in a way. So for example, imagine an area in the city that does not report a lot of crime, but has crime. And another part of the city that does. So usually you could just get biased and go to the area that has a lot of reported crime.

What can happen here is since you have this bunch of covarients, you would realize that there is another area similar to areas that have witnessed crime, and you would send police occasionally there to patrol. And you could actually discover crime. So that way you could actually get rid of some of the biases that exist in the data.

Eric Rice: This is a fascinating conversation and I think we need to move on to the final talk, but I've intentionally let the conversations between talks go a little bit longer because I feel like when you stack a bunch of them together then at the end, everyone remembers the last talk and everything in between gets lost and it's nice to have people respond as we go.

So Frank Dignum is going to be presenting Societal Challenges Need. Social Agents. Please.

Frank Dignum: Thanks. I should apologize [crosstalk 01:04:42]

Ayan Mukhopadhy: One moment. Talking about crimes, my laptop is missing.

Eric Rice: I stole it, I put it right over there, though.

Frank Dignum: So I think I misread the program I had 80 minutes plus two minutes of question, but it seems to be eight minutes and two minutes. So I will skip a few slides.

The second thing I want to apologize for. This is not done in cooperation with Melind. And it's the only one in this session. But maybe it's interesting anyway.

It does stand a little bit out from the rest. So this gives an overview of some of the projects that we are working on. And the type of problems that we look at has a lot to with the actual social interaction between people. So it's not so much people doing something and then trying to predict or doing some other thing with individuals, but more like what happens when they interact and how they influence and what happens by that interaction? And that happens in many different ways.

So one of the things that we did already quite some time ago is alternative currency, where we worked together with an NGO and in the Netherlands that does these alternative currencies all over the world and especially in developing countries.

And the question is, when does it work well, and when doesn't it work well? Because from economics they say, "Well actually if you look theoretically, it should work everywhere."

And it actually doesn't. So what are we missing? And it's a lot of this interaction between people. So we did simulations on that, and then you can actually change a lot of parameters and see what's the influence on that.

But a completely different type of application is this one, which is a communication training for students in medicine, and they have to learn to talk with their patients. And I don't know if you have any experience with talking with a doctor, a medical doctor. It's not very easy. They write even worse but they talk terrible.

So this is meant to train them. But if you want to train them you have to make realistic patients. So how does the patient react to the doctor? So the first version of this game was based on a kind of scripted dialogue where, it's very nicely done, where the professors could actually create the dialogues very easily. So they could actually indicate what students had to concentrate on. But the fact that it's all scripted means that students can kind of choose, have multiple choice and all the different points, which make it easy and not very realistic for them.

So in a normal conversation, you build a conversation, you have to actually watch and perceive what the patient picks up, how he reacts, and how you have to correspond to that. So we're now making a different type of game where you can have natural language input, which is far more free and where there's a more realistic conversation going on, which of course is quite difficult to do because it explodes the type of things that can be done.

So this is social simulation work, this is more like serious gaming, virtual characters work. And there's a lot of things that in between. There's this social sensing on demand, which is the crowdsourcing stuff within a town, getting people to send a lot of things, making sense of it, predicting what will happen. But then you have to incentivize people to actually send their data.

So we're doing this in the Netherlands where it works reasonably well, and they did a similar project in Jakarta, in Indonesia, and it doesn't work at all. And why it is, is because people don't trust the government. They say, "Well it was specifically for flooding. If we send our data to the government, they will see where things are worse and if it's not bad, terribly bad in a bad area, then we don't do nothing and we'll concentrate on the rich part of town."

So they rather don't send any data than having the government decide upon the real data.

So a question there is, how do people respond to your mechanisms? Refugee logistics is one that I'm just starting, it's about when refugees come in, which is quite current in Europe. How do we respond? How do we quickly the right materials and the right people at the right place?

Again here, what actually the government wanted from us was that we using AI techniques to predict. So they want a perfect prediction so that they know, "Next week there are 20 thousand refugees, the week after is only two thousand. And if we know that then we can respond to that perfectly."

I think you can use whatever AI techniques you want, but you're not going to predict this. So this is typically a case where we have to look, we're not just using our techniques, but we have to be in a dialogue with the government and say, "What do you actually want? Do you just want us to predict something for you and then that you can use it in your own organization to argue for more money? Or do you actually really want to help the people?"

So maybe we should look at the process. If you can't predict what's coming in exactly, you might want to change your process so that you can respond in different ways, make it more robust, more flexible. So we're going to build simulations where you can actually see the different consequences when you make different choices, which I think is far more useful than trying to predict something you can't.

Okay, given the time I'll move on. So just one remark on this. Who write the book "The Undoing Project?" Michael Lewis. No? It's a very nice book about the friendship between Daniel Kahneman and Tversky. And it has very nice first chapter on predicting the success of college basketball players in the NBA and how it was done first with simple rules by experts. Didn't actually work at all. Then they came in doing all his data-based thing, and it worked better, but they kept adding and adding and adding more and more vectors so it gets bigger and bigger and bigger, and then it didn't actually get better it got worse.

So how do we actually do this kind of interplay between experts and data? And when are rules getting too complex and you actually get in a lot of things that don't really matter but mess up the previous things? And when do they actually make a difference?

So sometimes people say it's going between simple rules and complex rules but it's: Sometimes simple rules can be very clever and they work perfect, or they can be very stupid. So it's about rhe difference between stupid and clever rules more than simple and complex. And that's the kind of things that in all of this we need to look at.

Okay this is the fisheries thing. We're doing a project, a European project on fishery management. And in fishery management of course, they have a lot of models for the ecological system. And the policy makers that want to protect the different species of fish.

So therefore they have all this complex ecological models, and then they see, "Based on this data, and this, the policy you can't fish more than this or you can only fish in that area or in that time of the year."

And then they're very surprised after one of two years, things didn't go as they planned. Why? Now they've come to realize that in between there are those fishermen. And they're not just variables that you can move with a policy, but actually have lives. They have economics, they have families that they have to provide for. They have to invest in their boats, they have to maintain their equipment, and there is a social system. So those fishermen, they live in a village and people depend on them and they live together.

So when one says, "Well I found a way to get out of this policy. I'll tell the rest." And they will all follow that, and suddenly the policy doesn't work. So what happens is that they come to realize that there is not just this one system that you have to simulate, but you have to actually simulate three systems with the connecting persons in the middle.

And this is something actually which is seem to be quite new, I was surprised to hear that. Of course they implemented people in the system, but the people were more or less like variables in the different system. In economical system they're rational agents. So they have their kind of parameters, utility function, and then they move.

In ecological system, they're capturing so much fish, or they do something and they have their own function there. Those functions are completely unrelated. And of course for people that doesn't work.

So what we're doing now is looking at, "How do you muddle this people that are in the midst of all this different systems that they're part of, in a way that is both transparent, that is explainable and that you actually see there is some coherence between all those different parts of a human life.

So some conjectures. People are social. I mean you open the door, you hear people talking. They're interacting. Social. So they have an inherent tendency to form groups and get together. Confine people, solitary confinement is terrible. It's one of the worst punishments you can give to people.

So we have to muddle somehow that people have to strive to actually interact with other people. It's not just a utility, it's a biological drive. And it shapes a lot of the ways that societies functioning.

They seek the familiar. So you keep doing what you already know because you know that that works, and therefore why would you change? You only change when it's really necessary. So that's why predictions often work. If you know something happened in the past you know people will repeat it again in the future. Not always, but by and large.

So if I make a policy, why would people change if there is one way of getting around your policy by not changing their habits? They will do it. So whenever people fave a way of getting around new things, they will do it.

One of the ways of getting this implemented is through social practices, and this is my new hobby horse of the last few years. There're everyday practices in the way they're typically and habitually performed in much of our society. So it's like going to work, cooking, meeting. All of that that come with a set of expectations of standard ways of interacting, standard ways of doing it.

These social practices, they work because everyone does them in a similar way. Not exactly the same, but similar ways. And therefore, if I start my presentation everyone is quiet. And sits and looks more or less. So if you would come into this room this morning and there were no chairs everyone would be very surprised.

"We would have a workshop here what happened to the chairs?"

So there's all those things that are standard, and that makes it very easy to do it in a certain way. But because of that, we also are shaping our way of doing things together based on this physical context and on the social context that we built up.

So a lot of things that we do, that we think we are very original and unique, we do because everyone does it in the same way. And this goes directly against this idea of that we are very rational doing everything with a utility. So we are doing it because we are always doing it. And maybe you can capture that in a utility but it's a different type of utility.

Let me skip this one. So what we need to muddle this kind of stuff is fundamentally socially defined individuals. We need something like norms, social facts, and all those kind of stuff as a normal part of reality. And we should have, in the end, cognitive abilities to create new concepts and categorize, classify reality.

Let me just give one example here which came up in Dutch politics a few years ago. We were in this crisis and then a politician started with, "We are now getting into the participation society."

Which meant that, "We're cutting the budged, we don't provide enough healthcare, and you have to help each other."

But coining this as a participation society suddenly made it feasible. People got angry, start discussing, but everyone talked about this participation society as if it was already there. So creating new concepts seemed to be very important for changing things.

Okay. Quickly on social agents. So what do we want more than things like [inaudible 01:19:57] or social motivations, identity, norms and habits, and practices? And motives, I took this one from McClelland, who argues this with a lot of research. We have an achievement motive that we actually want to reach a state, which are the typical goals that we know in agent world.

But also we have an affiliation motive. You want to have contact with other people, something we didn't model in any of our agent models before, which seems to be very important.

The power motive, which is not really the social power motive, power over another person, but the power to actually change the world. So the power to actually be able to do something, which if you have power over other people, you extend the things that you can do.

And finally, the avoidance motive. You don't want to get into trouble. So there are a lot of things that you want to avoid if you are a bit uncertain whether you can cope with it.

So those are the four things ... Yes, I'm trying to avoid any questions. So that's avoidance motive I was just getting there. So identity, it's not very simple. You have your personal identity, some professional, sportsman, family man. There are a lot of this different types of identities that you can have.

Yeah. I'm not in this picture, you don't have to look. I might be in this one though, I'm not in that one either.

So let me finish off with the social practice. Looking at this, everyone, at least Europeans will recognize this as a soccer match. And everyone knows exactly, "Oh, something is happening. There are two teams." People see this is the referee, there is the goalie, the goal. Players. So everything comes instantly.

The goal can be this one but can also be other ones and they still are recognized as goals, so there can be many instances, many ways of representing the concept of a goal. So we still have that all straight with this image and then there are things that we interpret right away. There's this off sideline, which this miracle line that somehow interferes with the playing. There's this fact that people with different shirts are opponents in this game, except if it's the referee. He is yellow. Everyone then knows that's the referee, probably because only one. Although the goalie might also have a yellow shirt.

So there are these things that we know instantly, we interpret. And also that come with roles. What people will do, what you can expect to happen and what they're capable of doing it. So the referee is supposed to know the rules of the game and know how to interfere with the game. All that kind of things that come instantly but are not determining completely what's happening, because every game is different. Every player has still a lot of room to make choices.

So we don't script the whole thing, though we give a kind of boundaries and context in which things happen. So it's a way in between the classical scripts and completely free, knowledge-based reasoning.

So we made a first [inaudible 01:23:48] where we say how would agents with this kind of stuff have to reason, and then you get this social practice line which looks at the context, chooses a social practice and go very quick to a line of action. If that doesn't work, if you don't recognize a social practice you can do a slow reasoning on utility based and look at your goals, make plans and all that.

So it's a kind of BDI still in this rightmost side, but a lot of things added that make it much quicker. So this is a fast and slow thinking of Kahneman and Tversky replicated here.

So conclusions already. There's actually one line of future work but there's quite a lot of future work. So the rest you can read yourself.

Eric Rice: Sorry there, so I know we're officially out of time but perhaps we can take a question or two unless people are anxious to get up and leave, in which case we'll violate some social norms and you can get up and try to leave.

Speaker 5: I actually did have a quick question. Can you go back to the last slide to the model? Yeah I see you have, so from salience you have context management, and then you have social practice and you have a little thing called revision over there, so I was wondering how you accomplish that in your model?

Frank Dignum: Well, yeah, that's a small question but a large answer. [crosstalk 01:25:18]

Speaker 5: I know. I work in the field so I know how difficult it is to answer that.

Frank Dignum: There are different revisions actually. So there's one evaluation so this kind of reinforcement learning, but also based on what you're experience is and what you recognize in the context, you can revise again what kind of things are really important in a certain context and what you will change.

And that might be that you make a different social practice for a sub category. So we have greeting and we have greeting with a handshake, for instance. And that's the kind of things that based on experience, both this one and that one, you will sub categorize social practice and maybe [inaudible 01:26:08].

Speaker 5: So this is a living, evolving [crosstalk 01:26:11]

Frank Dignum: Yep.

Speaker 5: Yeah.

Frank Dignum: So it's not a very simple thing to do. I also didn't say at all how we represent social practices, but I argue that you have to have two representations, one like a kind of neural network, very robust but very quick and [inaudible 01:26:33]. And then a kind of more knowledge based representation of that parallel that you use to explain and argue with.

So that's, the knowledge-based one also goes into this reasoning with the goals, but if you only do that one, you actually want a non symbolic representation.

Speaker 5: Thank you.

Eric Rice: I think we have time for one worm question.

Speaker 6: So you mentioned the importance of integrating like different kinds of models, so maybe you have a rational agent model for the economics, and like a OED or something for ecology and some kind of network model for the social element.

How do you do that? How do you put those things together? Are there general principles or do you have to ...

Frank Dignum: I hope there are general principles but we're still looking for that. So one of the things that I didn't really talk a lot about is this values. So what we want is a more fundamental model, that's why I got to this social agent or motives, they are far more fundamental than goals and all that.

Which gives you principles to analyze a situation, which gives you handles what kind of things of those different models you have to put together. More than that I don't think I can say in one minute.

Eric Rice: Fantastic. Well I think I'd love if we could all give a hand to all of the panelists who did a such a great job this afternoon.