

deeks

ASHLEY DEEKS: Thanks to you all for coming and joining me today. And thanks, of course, to Hannah and VJIL. I'm sorry Hannah had to duck out.

So my talk today is called High Tech International Law. And I've been thinking for a number of years now about the intersection between artificial intelligence and machine learning on the one hand, and international law on the other. So I work through, in a paper, how machine learning might affect things like detention by our military, how it could serve as potentially a tool for targeting decisions made during conflict.

And then, in a separate paper, I applied machine learning to the use ad bellum-- so questions that come up under the UN Charter and the international law related to the use of force-- where I was thinking about how machine learning might affect states' decisions about resorting to force, in the first instance. How it might be used, for example, to think through anticipatory self-defense or the ideas of proportionality in responding to an attack.

So I wanted then to step back and broaden the lens a little bit more to think about how machine-learning tools might shape international law more broadly. And I saw that there were scholars who were using these concepts-- these tools, these machine-learning tools, and computational text-analysis tools-- in a couple of areas, including in studying the Supreme Court and federal courts.

Just to give a quick definition here-- so computational text analysis is, basically, a process of using machines to analyze large corpora of data. So the idea is, you can use these tools to detect patterns in large buckets of data and identify the topics contained in those documents, maybe also the sentiments that they express. Are these documents in favor of something, or opposed to something? We also see-- and maybe you have seen-- reports about law firms and companies like Westlaw that are starting to use machine-learning tools to isolate particular sets of information. And I'll say a little bit more about that.

So I started to wonder whether international lawyers-- and, in particular, those in governments, those working for states-- could benefit from these kinds of tools. So I decided to try to sort through the combination of tools and problems that we're seeing in these other areas-- in the scholarship, academic area, as well as in the private law firm area-- to think about how international lawyers could use some of these same kinds of tools. So the paper is called "High Tech International Law," and it's coming out this spring in GW Law Review.

So I wanted to be creative in this project, but I didn't want to produce something that was well outside the range of possibility. I also thought that I needed to consider why it is the case that international law has been very slow to turn to these tools. What is it about international law that makes it conservative in this area?

And then I wanted to address a larger question too, which was, let's assume that states do start to use these tools, or some states start to use these tools. What is the distributive impact of the tools in terms of interstate power? Will it advance international law's goal to serve as a way to reduce

interstate tensions, reduce transaction costs between states? Or is it going to exacerbate some of those challenges?

We have seen unintended consequences in other areas of machine learning and AI. I'm sure you all have read about deepfakes. We've seen an overreliance on machine learning lead to embarrassing errors, such as when Amazon's program saw that you were buying one kind of fertilizer and said, maybe you'd like to buy this other kind of fertilizer. Others have used this to make bombs. [LAUGHS] So we need to think hard about these tools as we're thinking about adopting them.

We know that states have wide disparities already in technical competence and in an appetite for using technology in government. And yet it also struck me that this is a chance to lift all votes in some areas. It need not exacerbate the power dynamics, necessarily. And finally, as I'll talk about, I think there is a role for actors like the United Nations-- and also tech NGOs, who are interested in open data-- to step in and help develop some of these tools that maybe less technologically advanced states could use.

So my order of business-- I want to march through things as follows. First, I want to provide a little bit of background on machine learning and computational text analysis so we have a common vocabulary. Then I want to think about why states and international lawyers might choose to use some of these tools in their practice.

And then I want to turn to two buckets of possible applications. So one is in the development, the creation, of international law. Here I'm thinking about treaty negotiations and also discovering state practice for customary international law purposes.

So that's one bucket. And then the second bucket is thinking about how states could use these tools to resolve international disputes, so dispute settlement. And then I'll conclude with a couple of thoughts about the distributional effects that these tools might have in the real world.

OK, so by way of background, what are computational text analysis? What do I mean by machine learning? So painting with a broad brush, these are algorithmic tools that allow users to process large amounts of data quickly. They can extract patterns from the data that humans would not necessarily detect. They can make predictions, based on thousands of past examples, that the machine has processed.

And they can identify and categorize things. So you've seen you know algorithms identify dog photos versus cat photos. It can identify and categorize images, but also do that with documents, with words on a page or as between and among documents.

Computational text analysis is also sometimes referred to as text as data tools. Why? Because you're using the text, you're using the documents-- rather than, say, images or numbers-- as the underlying data that's being processed.

Machine-learning tools, which is a kind of subset of these, are tools that we think of as learning on their own. So you program and algorithm, make some prediction. You get a feedback about

how accurate it was. And then it goes back through, and it corrects. It reweights what it's evaluating to make the prediction to become more and more accurate.

So an example here is lung-cancer slides. So you might train an algorithm. You tell it, here are, 100-- let's say, here are 1,000 slides that show lung cancer. These are x-rays. It processes through the slides.

Then you show it another set of 100 slides, and it makes predictions about whether it's lung cancer or not. You then tell it whether it's gotten it right or not. It goes back through, and it becomes increasingly reliable.

Now, notice there, though, you won't necessarily know what the system's looking for-- what it is assessing in these x-rays-- that leads it to become more accurate. And so this is why you sometimes hear about these things as black boxes, because it may be hard to track exactly what the algorithm processed.

In computational text analysis, there are just a couple of concepts in here that I will refer to again, I'll describe here. First is this idea of topic modeling. So you're basically training the algorithm to sort documents based on whether they're mostly about international banking or they're mostly about climate change. But interestingly, you don't have to tell it in advance which buckets you think are going to be in the documents. You could give it 10,000 documents, and it might produce 40 categories of similar buckets of documents.

Another tool is something called sentiment analysis. And here, this is classifying documents, basically, as being in favor or opposed to certain propositions. So scholars have conducted sentiment analysis on Twitter feeds to try to anticipate the outcome of elections. Would that we knew the outcome of elections. [LAUGHS]

So private law firms are starting to use these tools to do things like identify the most relevant legal precedent, identify which arguments have been most persuasive to which set of judges. Maybe they can identify, how quickly does this judge tend to turn around or rule on summary judgment motions?

They've also used these tools to review contracts and, in some cases, even create templates for contracts. So they scan 2,000 contracts that the firm has worked on, and can almost take an average of what the basic contract is, and then compare it to the one that you've just developed for your client to see where the differences are.

And also, as I mentioned today, we have scholars who are working on domestic law issues who are using computational text analysis tools. So in the constitutional law space, we have people who are trying to predict the authorship of unsigned judicial opinions, based on stylistic analysis of each justice's opinion writing.

We have a number of law-review articles that have developed algorithms to predict the outcome of Supreme Court cases. And some of them claim they can do so better than our common law profs here in the building. And you can also see when courts are adopting language from the

party's briefs. It's basically a type of plagiarism software that can tell you something about which advocate you might want to hire. This is an advocate whose briefs tend to get adopted by the courts.

In administrative law, we're also starting to see this. So Professor Livermore here has done some work on this. And he's conducted a sentiment analysis of public comments on draft agency rules. So notice and comment procedure runs all the comments through to try to pick up what the thrust of the sentiments of the rules are.

OK, so that's just, by way of background, a little bit about these tools. Why might states be inspired to use machine learning or computational text analysis in their international relations or international legal settings? I think it is true that international law is generally very slow to change. This is the nature of the beast. International lawyers tend to be quite conservative-- small c conservative.

International laws form slowly. You all have studied this idea of customary international law. There is no such thing as immediate custom. It takes lots of time. Negotiating treaties takes lots of time. And we, necessarily, look to what we've done in the past. It is a pretty backward looking profession generally, and I think in international law, in particular.

Further, I think it's fair to say that governments often lag behind the private sector when it comes to technology. The government had a hard time, for example, getting people in to work on cyber stuff, because they're all busy in Silicon Valley making a lot of money. But I do think there's some reasons that states will adopt, or should consider adopting, some of these tools.

First, near-peer ambitions-- China, I think, is likely to dive in with both feet on this. So even if US or European lawyers are skeptical about the tools, they at the very least need to understand what the tools could be, even if they decide not to use them themselves. So I think pressure to understand and potentially to keep up with near peers.

Second, I think-- I've already gestured at this-- there's proof of concept in private law. Lawyers and private firms negotiate agreements. So do government lawyers working in international law. In private sector, you advise your client on the likely outcome of cases, you risk assessment. International lawyers have to do that with their policy clients, as well.

Conduct legal research on vast amounts of text-- both sets of lawyers have to do that. So we've seen firms start to use machine learning for all of these tasks. So there are these proofs of concept that international lawyers should pay attention to.

There might also be some client pressure to adopt the tools, especially, I think, if you have a political appointee who comes in from the private sector and says, I've just arrived. All my associates know how to use this stuff standing on their heads. Why are we still marching into the library with these dusty tomes of treaties in force? Let's step up our game a little bit.

To some extent, I think there will be-- the necessity of supplementing your ability to engage in your international legal work could help drive this, too. So I took a look at the number of lawyers

inside a range of foreign ministries, our equivalent of the State Department. The US has a huge number relative to everybody else.

I think we have about 250 lawyers in the State Department. Most other countries-- Mexico, Norway, Israel-- they have about 25, So a much, much smaller number of lawyers-- and yet, constantly, a new set of agreements, new sets of disputes, new sets of international case law that you're expected to get your arms around. Also, vast numbers of documents from places like the UN and travaux for the treaties-- decisions coming out of arbitrary tribunals, speeches being made by foreign ministries that are not your own, things you really want to keep up on-- so some sense that, we're going to need help getting our arms around all of those pockets of data.

Are there challenges to adopting some of these tools? Absolutely. I don't want to undersell that. I think one of the biggest challenges here is that the format that this data exists in is not what we would consider clean. So some of these documents are born digital-- since maybe, I don't know, 1994 at the UN. But everything earlier than that is basically scanned PDFs that have been uploaded into the library at the UN.

We also have vastly different kinds of documents that would be relevant. So in the private sector, you're mostly talking about court cases. Maybe you're talking about government regs, right? But they're all, more or less, in the same format.

Whereas, in international law, you care about a whole different kind of documents. You care about treaties, but you also care about speeches by foreign ministers. You care about what the Secretary General said in a statement in the Security Council. So you have that problem, as well.

And further, you have a language problem. So there are six languages in the UN, each of which is official. If you're an associate in a firm, maybe you do have to work with other languages. I don't want to suggest that all of the legal work is done only in English. But if you care about the state practice of 180, 190 states, that's a lot of languages you're going to have to get your arms around. So there will surely be challenges in moving forward on any of these projects.

Let me turn now just to give you some examples of how I'm thinking about these tools potentially being used in these two buckets of creating international law and then adjudicating. In the creation part, I've thought about it in two different areas. One is in treaty negotiations, and the other is in assessing the status of customary international law. So I'll start with treaties to just give you some examples of how I think computational text analysis might help an international lawyer do her job.

First, you could try to use these tools to identify your negotiating partner's equities. So for example, you could look at the frequency distribution of topics in every document that Kuwait has submitted in the UN for the past 20 years. Turns out, maybe you find that Kuwait always is talking about territorial integrity.

OK, so that's useful for you to know. When you're about to sit down with your Kuwaitian negotiator on some multilateral treaty, you know what buttons to push. You know what they're

going to care about in that negotiation, and also maybe outside that negotiation. Maybe there's a package deal to be had.

Or maybe you find out that India is fixated on disarmament. It's something you would never really be able to spend the time sorting through all those documents to find out on your own. But if you feed it to the right algorithm, really interesting things could come up.

You also could discover that groups of states tend to vote the same way on certain types of issues. Now, we all know that there are voting blocks in the UN. That's not a newsflash. There's something called the WEOG, Western European and Other Group.

It wouldn't be a surprise to find out that the Danes and the Norwegians tend to vote the same on lots of issues. But there might also be other stuff going on that you have not been attuned to that you could find out through these. And that would help you strategically manage the negotiation process.

You can even potentially predict the outcomes of treaties. So there's been some scholarship on this. There are a couple of scholars who have done this with Bilateral Investment Treaties, or BITs-- trained an algorithm on 1,600 BITs and created this algorithm that could predict the text of a BIT if you entered the two different countries that were notionally negotiating a new BIT, based on what the BITs looked like for those countries in other negotiations. And in particular, it applied it to a US/China negotiation. US is, in fact, in the process of negotiating a BIT with China, and it produced a draft text.

Another really common way that I think these kinds of tools will emerge-- and maybe already are being used-- is, during negotiations, interpretation and translation tools. So the UN has a formal translation and interpretation-- they have simultaneous interpretation among the six languages. But there are lots and lots of people there who barely speak one of those languages, and not their primary language. So tools like Google Translate and Skype's simultaneous interpretation are increasingly going to be important and helpful to those who are not fluent in one of those six languages.

Now, we've all seen some quirky slash silly Google Translate going on. But this is an area where the improvements are really going to be, I think, pretty steep and dramatic because these are tools that are being developed not just in private law firms. But every time we use it, it's basically learning and improving. And there's so many people using it that that's a tool that is going to get fine tuned over time. And Google has started-- and there are some competitors to Google, too-- using what's called deep learning, which are these deep neural nets, to help predict what word usually comes next after somebody says this word, what's the most common word that follows that, and so on.

Another way you could use some of these tools is you could web scrape your negotiating partner's media-- social media websites at home-- during the ongoing negotiations to see how the public is reacting to particular proposals. So maybe even in anticipation of your new extradition negotiation with Russia, you figure out, based on these tools, that the public really cares about the political offense exception. That would be good to know, as you're heading into the

negotiations, where your partner's sensitivities are based on where the public's sensitivities are, as well. Notice that that might work better for when you're negotiating with democracies than if you're negotiating with autocracies.

And then, more radically-- this is maybe one of the funkiest proposals in the paper or contemplated ideas in the paper-- you could run a motion-detection software during your negotiations to try to test your partner's truthfulness. So let's say you and the Chinese are having a negotiation about those BITs. And you're in your home conference room, doing it. Turns out you have cameras embedded in the ceiling that are recording the conversation and also running a motion-detection software over that video.

And on a break, you go back into your little room. And your experts say, OK, when they said that they could not-- this was their bottom line on this-- they were lying. So that is another place where you could use these tools.

That's treaties. How about custom? How would this potentially be relevant for customary international law? So I think, as most of you know, state practice takes a huge range of forms-- diplomatic acts, legislative acts, submissions to the UN, public speeches. And one of the criticisms of customary international law is that it's very often driven by Western states.

And in particular, historically, states like the US and the UK published these digests of their practice and, every year, put out huge volumes. So easy place to go to try to find out, well, what has the state practice been on X? You have your [INAUDIBLE]. But there are many, many other states engaged in all sorts of state practice, now and historically, where that practice has been much harder to detect. So clearly, the trick here is identifying either past state practice, or even current state practice.

And here, topic modeling could do a great job. If you give it bucket of documents, it could help extract practice on territorial seas or on overflight rights. Let's say you decide you're interested in freedom of navigation. Not only could you topic model for that, but you can also-- it will find things related to that, even if you don't tell it what other terms are commonly associated with it. So could determine that related concept include disputed territorial seas, or is this a rock or an island-- things associated with Law-of-the-Sea disputes.

So the long pole in the tent here, I think, is digitization. It's not the machine-learning parts of this. I think the machine-learning tools could, relatively easily, be developed. It's getting the stuff into a bucket where you can process it. And even if these archives prove too hard or too costly to unearth, you can web scrape, I think, the practice that's captured online in a way that gets you beyond just a Google search. I think it goes deeper than that.

OK, so that's potentially creation of international law. What about international disputes? Can the tools be useful here, too? So I think, potentially, yes. I think you could use computational text analysis and machine-learning tools to better prepare for litigation, and also litigation in tribunals, but also dispute resolution outside of tribunals in a more traditional diplomatic negotiation shuttle diplomacy between two states.

On the first part, on the tribunals part, I do think these tools can help you understand the workings of arbitral role and judicial tribunals. You could figure out, for example, which arbitrator wrote the underlying arbitral opinion. Usually, there are three arbitrators. And usually, the opinions are not signed. But you could conduct a stylistic analysis to figure out who wrote it.

You can also help figure out which arbitrators have more or less influence in the drafting. And that would tell you, as a state, which arbitrator you might want to hire or not hire in the future. Some people have figured out that sometimes it's the secretariat of the arbitral tribunal, the staff that's running the arbitration, that are really behind that the outcome, not the actual arbitrators. So that's also interesting to know.

I mentioned plagiarism-detection software could be helpful to know which judges or which courts borrow from party's briefs. Why? Well, it would inform whether you hire outside counsel expertise or you just go with your own lawyers. Could also be relevant in your decision about whether you appear before a tribunal at all.

So there have been a couple of instances recently, and some historically, where a player decides not to appear in a court case or an arbitration. Notably, US decides not to appear in the Nicaragua case and ICJ after they lose on the jurisdictional front. More recently, you have China deciding not to participate in an ITLOS hearing. If you know that the tribunal relies heavily on the party's briefs, maybe you make a different decision on that. So that could be an option, as well.

And then, how about predicting outcomes? So I think this tool seems among the more speculative that I talk about in the paper, although people have done it, as I mentioned, with the Supreme Court. And other scholars have tried to do it with the European Court of Human Rights, which produces thousands of judgments-- big corpus of work. And they've claimed that they had about 75% accuracy.

So what they did is they trained it on the facts, and then the judgment. Then they would present the system with just the facts and see how it did at predicting the judgment. So notice one thing here that makes this an imperfect study is that the people writing the facts might well frame the facts in a way that supports the judgment. So the facts are not necessarily neutral facts. But several scholars have actually done work on this area.

So notice it would be hard with something like the International Court of Justice, where there are very few decisions in any given year-- maybe 175 total, in the history of the court. That's not a big enough in to do. You should just read those opinions. This comes in handy where you have a large volume of stuff that it's really hard for you to read yourself.

How about arbitrating algorithmically? So this is something that a couple of scholars have proposed. I think, more likely, what you could use the tools for is to sift through the law more quickly and more effectively than a human might be able to propose settlement ranges or, potentially, draft awards. You can have the system draft an award, and then the humans could compare their draft award to the system's award to figure out whether they track or not, and if not, why?

How about for disputes that are not inside a courtroom or inside an arbitral room? What about just our straightforward policy disputes? So Middle East peace process is sort of exhibit A, but there are a whole lot of these persistent foreign-policy disputes. So take Western Sahara, where Morocco claims it. Western Sahara wants to be independent. You could use these tools-- web scraping, for example-- to make sure that you have unearthed all possible ideas that have ever been proposed about this topic.

So let's say I'm the arbitrator, or I'm the diplomat selected for adjudicating the Western Sahara dispute. I'm finally going to get it right this time. Well, one thing I'd like to do is know, what are the 300 proposals that smart people who've been thinking about this for a long time have developed? So this could be a way to help get that in front of you.

And then, finally, I've mentioned the idea of understanding public reactions of affected nationals. That's relevant, I think, both before negotiation and after, or during a negotiation. But notice that these tools can also be used to manipulate public reactions. There is a dark side here that we should keep in mind.

So you could develop tools that could sway the public toward or away from certain proposals. And you could even do this as a third-party state. I, as the arbitrator or the diplomat, might not-- I want there to be an honest broker. But there could be some state out there that wants to put a [? spanner ?] in the works. And they could be doing all sorts of things in this space to try to throw my efforts off.

OK, so those are just some ideas to chew on for ways in which these tools might be used. Are there challenges? What are some of the ethical challenges here to deploying these? So if you guys have read any news articles about machine learning and AI, there are challenges in getting this right. I'll mention just a couple.

Many of the same problems that you've already read about in the Amazon's recommending bomb-making materials articles, I think, would come in play here. So the first is embedding or replicating bias. Let's assume that arbitrators tend, in general, to find for wealthier states because they're worried that the wealthy states could really make a big fuss-- bad mouth them publicly in a very broadly circulated way-- if they hold against the wealthier states.

So if you're training an algorithm that's trying to predict the future outcome of arbitral disputes, it's going to take into account the fact that wealth matters, even if you don't tell it that wealth matters. If you enter in information about the countries involved, it will ascertain that that's an underlying factor. It might make the prediction more accurate, but it might be a bias that we don't want to replicate.

How about errors in data? So in all of this, it's garbage in, garbage out. Your predictions are only as good as the quality of the data that you have. So it'll take a lot of work to get this right.

Opacity is another critique, this idea of a black box. I think if you're a lawyer confronted with some of these tools-- somebody says to you, hey, policymaker X is really hot on doing topic

modeling in this set of documents-- you want to know, OK, what went into this algorithm? Why should I trust it? Can you tell me a little bit of how it works?

And then, finally, there's this idea of automation bias. So people tend to believe what machines tell them. And so if your machine is wrong, but seems credible, you may go with it, even if your own experience is to the contrary. So that would be another thing you'd want to get trained on before you dive all in on this.

OK, so distributional effects-- if states do adopt some or all of these tools, who are the winners, and who are the losers? I think this is an important question to think about. And I think the story here is mixed. If we're talking about tools that are relatively inexpensive to develop, relatively easy to use, then I think those are the kinds of tools that are going to lift all boats. So if we're thinking about the translation or interpretation software, for example, that helps improve the capacity of states that are less technologically savvy to participate during multilateral negotiations.

What about the tools that are sophisticated and costly to develop? Well, will those exacerbate existing power differentials? Quite likely, yes. Tools that help predict the outcome of court cases-- if we can get that right. Or the roles and power of different arbitrators-- same thing. Motion detection software deployed in secret-- I think, clearly, not all states will come to negotiations with those tools in their pockets.

Of course, it is true that states already come to negotiations with very different power dynamics. The question really is, how much is this altering that? Is it doing so in a disproportionate way?

Just as a footnote, I think there will also be distributional effects within a single state, so not just between states. But these kinds of tools, if they start to come into government work, will empower those in government who are tech literate. They will empower lawyers who are comfortable dealing with technology. And they will disempower the Luddites, if these turn out to be valuable tools.

Overall, I could imagine four different outcomes here. So one is we say, look, these tools are overall going to be beneficial for international law because states will be more easily able to come to agreement in negotiations or in adjudication. Option two would be, well, yes, they'll facilitate agreement. But they'll do so in a way that seems unfair to those who are not able to use these tools. So it actually exacerbates dissatisfaction with international law.

Third, if they're overtly seen as misused or used unfairly, it might actually hinder agreement. States will not want to conclude agreements where they feel like they've been buffaloed. And then, fourth option is, no notable changes, either because the tools are ineffective, or because everybody's got them. And so it raises all boats in the same way.

And it may well depend on the context of the interactions, the context in which the tools are used. So if two states are in a zero-sum negotiation, they're not very likely to tell the other side, look, we've got these tools. There they're running on your past documents. And we know things you don't.

But if you're in a negotiation where you're really trying to get to a win-win treaty, you might be more willing to tell the state, look, we ran a study of all of your past treaties. And we found that, if we conclude this, it's going to be in tension with two of those treaties. So let's talk about that. So you could do it in a positive way, too. So I think, most importantly, is how states use the tools-- how do they use them, not whether they use them.

And I'll just close by mentioning this role for the UN. I think there is, potentially, a role for the UN secretariat here. So it could do a couple of things. It could say, look, we're going to clean the data. We're going to create this giant database, starting with what we have, and maybe cleaning up more historical documents, and make it readily available for researchers.

I think they could also offer a data-scientist-type services, including some of the algorithms that I just discussed. They could develop some of these and allow states to play around with them, put different selections of documents in, and produce different sets of information.

And there are a bunch of publicly-minded data NGOs. They're interested in public data open-access work. They've done a lot of work with US cities. Those kinds of actors might potentially, if they're interested in foreign policy, find an opening here for new, interesting work.

So just to conclude, I think there are lots of room for actors who are interested in applying machine learning and computational text-analysis tools to get creative in the international-law space and the foreign-policy space. And so I would encourage those interested in international law-- like I assume you, in this room-- to keep an eye out for these tools and think about, when you see them, is this something that could help advance the progress of international law?

So with that, I will wrap up and see if you guys have any questions. Please. Yes.

AUDIENCE: Sorry. You mentioned that this machine-learning software could be an equalizing force that raises all boats, or it can be used particularly by more technologically-advanced countries. Is there one side, just from general trends of advancing technology, that you see that way more towards? Do you think the UN will be able to propagate some software like this throughout all the bodies? Or will it be the US and other advanced countries?

ASHLEY DEEKS: So I guess I would say it really, again, depends on which tools we're talking about. So right now, if you're a negotiator from Botswana, and your first language is not one of the six UN languages, and you come to the negotiations, you might already be using Google Translate or apps that help you readily translate your proposals into one of the six languages. That may already be in place for some of those tools. But more broadly, I think it's the case that there are a handful of states-- maybe you can list them on two hands-- that are really interested in AI and machine-learning tools, have private sectors that are really interested in them, and are making significant advances.

So Chatham House-- which is one of the few actors that's thought about this in the international space. One of their proposals is that states really invest in these tools-- not just China, and Russia, and the US invest in these tools, but all states invest in these tools-- to try to avoid falling

behind. And so I think some of it will just depend on the appetite of different leaders to pick up that call.

Same thing with the UN-- so the UN is involved in just finishing up a study about customary international law. The ILC-- the International Law Commission-- and Michael Wood have done a long-term project on this. And one of the things they've been wrestling with is, how do we find customary practice? There are all these documents out there. How do we do it, urging states to come forward with their own practice?

But they don't talk about the potential for machine learning and computational text analysis here. And I feel like that's an opportunity missed for them. So if somebody from the UN Secretariat said, hey, could we talk? What do you have to suggest? I think there is a lot of room here for them to do good. So we'll see. We'll see. Yeah.

AUDIENCE: So I feel like the two problems here of the difficulty in programming in for the inputs and principles on the front end, and then the difficulty in figuring out what the algorithms used in making the recommendations from the back end, are more acute in the due-process context of making decisions to target in the field or making sentencing decisions in court cases.

ASHLEY DEEKS: Yeah

AUDIENCE: But in the predictive cases that you're talking about, couldn't there be the same problem where-- I feel like there's a lot of head faking and politicized language in travaux for the treaty negotiations in UN Security Council. Are you worried that algorithms might spit out bad recommendations that people assume, this is all about sovereignty, because China talks about sovereignty all the time? People might be able to sort of manipulate the algorithms, or at least that they'll spit out recommendations that are not very helpful.

ASHLEY DEEKS: Well, so where I thought you were going with that question is about gaming the algorithm. So I think part of your question may be-- OK, let's say that China knows we're running algorithms on all the documents that we have, all the documents the UN's got, and that that's going to help effect our negotiation posture on things. And so China says, OK, we are going to make lots of statements about the importance of Botswana. And so our algorithm says, oh, in the past five years, China's been obsessed with Botswana.

But it turns out it's just a head fake, right? And so when we say, well, we're happy to provide some more financial aid to Botswana if, whatever. They're like, oh, OK, I guess then we can agree on the deal, because they've totally set us up to head down the wrong path. I think that's absolutely a possibility. But I think your question had more in it, too.

AUDIENCE: Well, I think my worry there is that, because of that effect, that you might get bad algorithmic recommendations and not be able to figure out in the same way-- I worry that the algorithms would not be able to understand China is gaming us. Or even not gaming, but I feel like there's a lot of language in resolutions that tends to say what one side wants it to say, but be said in noncommittal resolves-- feel strongly that, and not decisional language.

I mean, I guess you could just program around those problems. But I just worry that, as long as that this black-box problem of not being able to understand what inputs mattered to the decision exists, that it'll be hard to rely on the recommendations that come out of them. Even if it's not the algorithm making a decision, even if it's just offering you preparatory advice that, as long as you don't know what the logic behind it is, that it won't really be very useful.

ASHLEY DEEKS: OK. So there, I think the answer is explainable AI. And it could well be that the kinds of algorithms we're talking about here don't require deep neural nets that are really hard to actually make explainable. That these are things where-- so let me back up.

Explainable AI is this idea-- it's trying to wrestle with the black-box problem. And there are a whole bunch of different types of explainable AI that computer scientists have been developing. Some of them try to crack open the black box. Some of them write algorithms alongside the main algorithm and give a rough approximation of what weight the main algorithm has given to different factors, without being a perfect representative. So I think that's the direction you would want to take these.

So A, I don't know that you need the deep, deep neural nets-- except for things like translation-- for the kinds of algorithms that I have in mind. And then, second, it would therefore be easier to create explainable AI for the lawyers that are trying to use this, or the diplomats that are interested in using this. Any other questions?

AUDIENCE: So I was thinking just-- I don't really understand, and I think this is and partially because I'm not very technologically savvy. But if countries in their own language use-- I guess it's just I don't understand how a computer would be able to, in one language or multiple-- just the difficulties in translation of different languages, and how there are some terms in other languages that don't translate into what they should mean in other languages. And I just see a lot of problems with that.

ASHLEY DEEKS: Well, so I think if you were having a US computer scientist try to create algorithms that were sensitive to a Nepalese dialect, [LAUGHS] that would be a problem. You would have to spend the time getting it right on the front end. And maybe that's you get a high-level Nepalese interpreter sitting next to you saying, this word doesn't mean what you think it means. This is going to pick up this set of documents, and so on. So absolutely, you'd have to be very careful about that.

But the six UN languages-- computer scientists are training their algorithms on that because it's such high-quality translation across six pretty different languages. So that would be the gold standard. But you're absolutely right.

I mean, that is a real challenge. And especially if you're in my customary international law state practice project, trying to figure out what phrase Vietnam uses for rock versus island will be really important. So it's a hurdle, for sure. And more of a hurdle-- I think, that's why domestic law doesn't quite wrestle with those challenges.

Well, thank you, guys, very much for coming. I appreciate the time.

[APPLAUSE]