

Conversation 2

Dispute Resolution Transformed by Technology

ABSTRACT

This conversation explores how technology changes the way disputes are solved. The focus is on the impact of artificial intelligence. After reporting on a competition, in which lawyers and an artificial intelligence competed to accurately predict the outcome of disputes before the UK Financial Ombudsman, the speaker explains how artificial intelligence is practically used in dispute resolution. Such use cases include the production of information, the creation of focused analyses, the finding of decisions and the generation of communication. The speaker then presents research projects using artificial intelligence to predict dispute outcomes in the courts of different countries. The conversation also addresses the ethical questions arising from different use cases of artificial intelligence in conflict resolution. In conclusion, the potential of artificial intelligence to improve access to justice is identified together with the ethical challenges that need to be addressed.

<i>Speaker</i>	Felix Steffek
<i>Moderator</i>	Mihoko Sumida
<i>Comments</i>	Kazuhiko Yamamoto, Masato Fujita, Keisuke Takeshita and Hiroaki Yamada
<i>Concluding Conversation</i>	Felix Steffek and Mihoko Sumida
<i>Questions for Further Thought</i>	Felix Steffek

AI VERSUS PROFESSIONAL LAWYERS

Sumida: The second session will be on ‘Dispute Resolution Transformed by Technology’. I would like to ask Professor Felix Steffek to begin his presentation following his self-introduction. Thank you very much.

Steffek: It is a great pleasure to be here with you today. I am very grateful to Professor Mihoko Sumida for arranging this amazing course. It has been a great pleasure working with you, Mihoko. I am also very grateful for the excellent translation provided and for the colleagues and external visitors to be here. It is great seeing so many of you in this virtual setting.

My name is Felix Steffek, as Professor Sumida has already mentioned. I am based at the University of Cambridge, Faculty of Law, and I am the Director of the Centre for Corporate and Commercial Law. My interests are in dispute resolution, but also in corporate finance and corporate insolvency. Most recently, I have engaged with LawTech, also referred to as legal technology, particularly with artificial intelligence (AI) in dispute resolution. The latter is also the topic of my contribution to this session. In this contribution, I will focus on civil and commercial dispute resolution. I will not cover criminal law and particular public law issues.

How do I use the term 'AI' in this lecture? There are many definitions, but a very brief and rough-and-ready definition would be technology trying to make intelligent choices. When I use the term 'AI', I have in mind machine learning, in particular deep learning, and natural language processing (NLP). Later, I will give you an example of how it works, and I hope that it will also help you understand what AI in law means.

Here is a brief outline of what I will talk about. To start with, I will give you a report on the Case Crunch Lawyer Challenge. It was a competition that students from the University of Cambridge, Faculty of Law, had staged. It was really interesting, so I thought I should tell you about it. Then, I will provide some background information on this competition. I will explain how the students used AI in the area of law. Following this, I will report on current research. I will refer to studies on the use of AI in dispute resolution so that we get an understanding of the field and where the research stands at the moment. Then, I will move on to discuss some use cases. In other words, I will talk about how AI can be employed in the practice of law. Then, I will discuss three essential issues. *First*, I will address the issue of bias that concerns many people as regards the use of AI; i.e. AI perhaps makes the same mistakes as we have made in the past. *Secondly*, I will deal with the issue of transparency because law and legal decisions have to do with transparency. For example, in court, we do not only hear the result but also want to know why a judgment is made. *Finally*, I will bring these issues together to raise more general ethical considerations. I will investigate the ethical advantages and disadvantages of using AI in dispute resolution.

Let us start. What you see in Figure 2.1 is some key information on the Case Crunch Lawyer Challenge. As I mentioned earlier, I will report on this project that was conceived by students at the University of Cambridge. What happened? In October 2017, they organised a competition between an AI and real lawyers. You can see that this was already some years ago. Later, I will talk about the advances in AI and law that have happened since then. At the time, these students were really at the forefront of developments. What was this project about? There were four




FIGURE 2.1 Lawyers vs Case Crunch technology

students. They had programmed an AI to predict the outcome of cases. Their idea was to test how good their AI was by asking lawyers to compete against it. Both the AI and the commercial lawyers would predict the outcome of cases. The challenge was to predict more cases correctly. In the end, more than 100 lawyers participated. They were all very good commercial lawyers, mainly based in London.

What was the specific task? The task was to predict real cases that the Financial Ombudsman Service had to decide in the past. The Financial Ombudsman Service is an ombudsman service for the financial services industry. At the time, there were many cases concerning Payment Protection Insurance (PPI). This is an insurance product that enables consumers to insure repayment of credit if the consumers, i.e. the borrowers, become ill or disabled, lose their job or face other circumstances that may prevent them from earning income to service the debt. In the past years, such PPI was widely sold by banks and other credit providers as an add-on to a loan or an overdraft. However, consumers often felt that they did not need this insurance and that banks or financial institutions had pushed them to buy this product. Later, they wanted to get rid of it, saying that they were mis-sold such PPI.

Everyone taking part in the competition, the AI coded by the students and the more than 100 commercial lawyers, received real cases from the Financial Ombudsman Service. Of course, they only got the facts. They did not get the outcomes, i.e. the decisions that the Financial Ombudsman Service had taken. Based on these facts that describe the case – how the consumer bought the product, what the consumer was told and so on – both the AI and the lawyers had to predict whether the consumer applying to the Financial Ombudsman would win or lose the case.

The facts of thousands of cases were put on a website, and the lawyers could log in and look at these facts. Then, they could take all the time they wanted and could consult all the material they wanted, including both online and offline sources.

Then, they had to click a button to predict whether the consumer would win or lose the case. In other words, the lawyers had to predict the decision of the Financial Ombudsman Service. This way, the lawyers made more than 750 predictions. So, the lawyers were quite active, and there were a large number of predictions by them. The AI did the same. After all predictions were made, everyone attended a conference in London to learn about the results. To make sure that this was a fair competition, there was a technical judge who inspected that everything was fair in terms of technology and that no one had an advantage. There was also a legal judge to make sure that the competition was fair in terms of legal matters. I had the privilege to serve as the legal judge. I can assure you that it was a fair competition. No one had an unfair advantage over the other. In Figure 2.1, you can see the result of the competition.

The lawyers achieved a prediction accuracy of 62 per cent. That means that in ten cases, they predicted about six correctly and four wrongly. The AI, the students called it Case Crunch Technology, had a prediction accuracy of 87 per cent. This is close to 90 per cent, which roughly means that in ten cases, the AI was correct nine times and got one case wrong. Hence, the competition resulted in a win for the AI. Put differently, the AI was better in predicting the outcome of PPI cases than the lawyers.

After revealing the results, the participants at the conference in London also had a discussion on what these results meant. In the discussion, the issues that were raised concerned costs and time. For each case, a lawyer would probably take one hour (at least) to work on it. Quite likely, there would also be other staff in the law firm working on the case, which would probably result in a three-figure fee to be paid by the consumer in British pounds. On the other hand, the AI, once it is set up, which comes with costs, would do the calculation and the prediction in a fraction of a second. Depending, of course, on the number of cases, the AI would incur fewer costs than the lawyer's firm would charge.

So, it is worth noting that the discussion at the conference did not only concern the accuracy of the prediction but also the costs and the time involved. Of course, the main issue at the conference was the result of the competition, and this concerned the accuracy of the prediction. The discussion at the conference also concerned what this might mean for the future. A further question that was raised concerned the amount of time one needs to invest to get information on the law and what the costs of this are. This concerns both the AI and the lawyers – both need to acquire information on the law before they can render their services.

We were quite happy to see that the results of the competition were reported all around the world. On the right side of Figure 2.1, you can see a selection of journals and news providers that reported on this competition. I will come back to that later and tell you a little bit more about what the British Broadcasting Corporation (BBC), said on their website. I think that the wide media coverage and interest show that people are interested in the future of justice.

WHO WOULD YOU LIKE TO ASK FOR LEGAL ADVICE, A HUMAN
LAWYER OR AN AI?

Steffek: As Professor Sumida intends this event to take the form of a dialogue and I very much like this idea, I have a question for you. Imagine you had a dispute concerning PPI, as was the case in this competition, with your bank and you considered an application to the Financial Ombudsman Service. However, you would be interested in knowing whether it is worth going to the Financial Ombudsman Service, investing your time and perhaps your money. You want to know whether you would be successful. Who would you ask?

I will provide you with three choices, but you can also make up more choices if you have another opinion. Would you ask: (A) a human lawyer, (B) the Case Crunch AI or (C) a human lawyer who uses the Case Crunch AI?

Student A: I would like to consult both, a human lawyer and the AI, so I would choose option C. The reason for this is that I have concerns about the diagnostic results of AI alone. I am not yet familiar with AI, and I would like to know information surrounding the results, such as why these diagnostic results were obtained and what rulings have been made in the past. Such information is not available in the AI's diagnosis. That being said, I just saw that the AI had a much higher percentage of correct answers, so I would like to benefit from both.

Student B: I would choose only the AI, option B. The reason is the cost hurdle. There are enormous legal fees involved in asking a human to give their advice. Accordingly, I would be happy with the AI alone, as long as it is accurate, and I would be able to rely on it to prepare my decisions.

Sumida: Are there any other voices with different opinions? Would anyone opt for option A, i.e. only consult a human lawyer?

Student C: I would prefer to ask a human lawyer, option A. I learned from a professor at my university that law has a social function and intervenes to support people who tend to become emotionally involved in a dispute. Such intervention can come in the form of a lawyer, a third party. The lawyer will aim to reach a social solution while also calming the emotions of the respective parties. I agree with that very much. When you are faced with the need to use insurance due to a contingency, such as loss of employment, you may have a lot of emotions and may not be satisfied with the situation. As a party, you may feel uneasy about whether you can make a good enough claim by pitching it to AI. Instead, I feel that if I ask a trusted lawyer, they will be able to adequately get my point across. This is the first reason for choosing option A.

Secondly, I do not know from what perspective the accuracy of the AI's prediction is measured. There might, perhaps, be aspects that are not correct in the assessment by the AI even if the assessment is in your favour. For instance, an AI might suggest that you are entitled to claim insurance benefits, even though in reality you have a very weak case. Hence, I think that asking a human lawyer is not necessarily a negative course of action.

Sumida: So, we have had a lot of opinions. Felix-san, would you like to respond?

Steffek: Thank you very much for your opinions. These were perfect answers because you covered the relevant issues. Often, we concentrate on accuracy first; we look at the facts and want to predict the future. In addition, you mentioned the issue of costs. This shows that while we desire perfect justice and predictions as an ideal starting point, in practice, we ask the question what this will cost. Sometimes, we will accept a less-than-perfect service because the costs of a perfect service would not be justified for the specific case at hand. This applies, in particular, to cases with a low value. There are also further considerations. The last comment emphasised the human side of justice. Justice concerns not only cold facts but also emotions. Taking all answers together, I think you have covered the relevant issues.

Interestingly, the fact that we had voices for all three options shows us that different people want different things in justice. There may be people who are only interested in the result. They say, 'I want the money back and that is all.' For other people, dispute resolution is rather an emotional process. They have anxieties or they are angry; they want to express that in the process of solving their dispute, and they want to express it to a human. Different people want different things in dispute resolution, and it is for them to decide. This is why your answers are all good answers. There is no wrong answer because it is your wish. Different citizens have different wishes. This reminds me of my little son, Max. He is four years old, and he came back yesterday from kindergarten and said, 'Daddy, people are different and they want different things.' And that is how it is.

However, we also have constraints. Such constraints set limits to our wishes. We only have a certain number of resources, and when you go to court, it can be a zero-sum game, i.e. there can only be one winner and the other party will be the loser. In reality, we cannot always have it our way. This is why we need the law and lawmakers to make choices, which wishes to grant and which wishes to deny. This is the background of what we are discussing today.

CREATING LEGAL PREDICTION TOOLS

Steffek: Now, I would like to explain how the team behind the Case Crunch Challenge organised their project and how they used AI. It all started with a dataset. They went to the Financial Ombudsman Service and asked whether they could have a dataset of cases that were decided in the past. The Financial Ombudsman Service decides hundreds of thousands of cases every year and provided the team with a large dataset. The Case Crunch team decided to concentrate on cases concerning PPI and not deal with other cases. Each case had the same structure. The first part provided the facts, and the second part the decision of the Financial Ombudsman Service. As there are many PPI cases in the United Kingdom, the dataset was quite big. There were around 100,000 cases in the dataset. The team then took 10 per cent of these cases, i.e. 10,000 cases, and created a training dataset with these cases. They used NLP to classify these 10,000 cases as cases where either the consumer wins or loses. To give

you an example, the team used an algorithm according to which if the last sentence of the decision contained the phrase ‘not successful’ or the word ‘unsuccessful’, then that case was classified as ‘consumer loses’. This shows that already at the stage of creating the training dataset, they used computer algorithms. This was necessary because it would have been too tedious to classify all 10,000 cases by hand.

You can imagine the result of this algorithmic classification as a big table where in one column, i.e. in one cell for each case, you have all the facts, and in the second column, you have a ‘yes’ or ‘no’ as regards the outcome. ‘Yes’ means the consumer wins, and ‘no’ means the consumer loses. These two pieces of information, facts and outcome, are linked to the relevant case, identified by a case number. After the training dataset was constructed in this way, the Case Crunch team used a neural network, which is a form of AI, to learn which factors in the facts would determine the outcome, i.e. whether the consumer’s application to the Financial Ombudsman Service would be successful or unsuccessful.

Thus, this is where AI comes in. The neural network is an algorithm that trains itself. You do not have to do anything except wait until the training process is concluded. Essentially, one provides the algorithm with the relevant information – in this case, the facts and the classification as ‘win’ or ‘lose’ – and the neural network will then, by trial and error, determine which factors within the facts determine the outcome. During the training process, the AI makes predictions. These predictions will be very arbitrary at the beginning. It will start with the facts of one case and arbitrarily predict the outcome. Let us assume that the consumer, in fact, won this case. If the AI predicts a win, then it sees the factors it considers as confirmed. If, however, the AI predicts a loss and then realises that the prediction is wrong, it will go back and try to understand what factors are relevant for losing. The AI does this process again and again for all 10,000 cases in the training dataset. The higher the number of cases the AI can learn from, the better the predictions can get.

At the end of this learning process, the AI gets better and better through trial and error and can connect factors in the facts with the expected outcomes. This process of learning which factors are relevant (and which are not) is not something that the human influences at that stage. Instead, the neural network, of which you see a representation here in Figure 2.2, does it independently. The input layer in the figure is the facts, the output layer is win or lose, and the circles in between are the factors which the AI determines to be contributing to whether a case will be won or lost.

The relevant factors are determined and weighed by the AI itself. It is not the human creator who does that. It is the neural network that does it. Notably, since it is the AI determining the factors, the human creators and users do not always know what the AI uses as factors to determine whether the consumer will win or lose. This is called the ‘black box’ effect. Due to this effect, humans have problems understanding which factors are relevant for the AI, particularly since the AI looks at a large number of factors. Additionally, as shown in the figure, the connections between the factors are quite complex. They are often so complex that it is difficult for humans to fully understand the weight of each individual factor for the final result.




FIGURE 2.2 Artificial intelligence at work

At the end of the training, the AI is ready to look at new cases. This is when the Case Crunch competition started. After the AI had been trained on 10,000 cases, it would then – in the competition – get the facts of a new case and predict the outcome, i.e. whether the consumer would win or lose. At this point, the AI had a prediction accuracy of 87 per cent as I reported earlier. In other words, it would reach the correct prediction in 87 per cent of the cases. This is an example of how AI can be used to predict the outcome of a dispute resolution procedure such as ombudsman proceedings and court litigation.

LEGAL AI: WHERE DO WE STAND?

Steffek: Let me provide you with an overview of some of the existing research and literature on using AI to predict the outcome of cases. In these projects, the fundamental approach is often the same, mirroring the methods that I have just explained. Researchers take a dataset, train an AI and then apply the trained algorithm to a new case where the AI does not have the result but will predict the outcome. The result is then compared to what the judges or the ombudspersons decided, and the precision of the prediction is calculated. There are more studies than I can present today. I picked some from around the world to give you an impression of where research currently stands in this field.

I will start with the paper by Katz et al. published in 2017.¹ They looked at the US Supreme Court. This research project had 28,000 decisions in the dataset. An important question concerning such datasets is: ‘What exactly is the information contained in a dataset?’ In Figure 2.3, I refer to this as the basis of the research project. This particular research project only had meta-factors in the dataset. What are meta-factors? ‘Meta-factors’ refers to abstract information on the case and excludes the specific facts of the case. For example, this research project used the following meta-factors: *First*, what is the circuit of origin of the case? In other words,

¹ Daniel Martin Katz, Michael J. Bommarito II and Josh Blackman, ‘A General Approach for Predicting the Behavior of the Supreme Court of the United States’ (2017) 12(4) PLoS ONE.

Authors	Year	Dispute body	Basis	Decisions in dataset (rounded)	Precision (rounded)
Katz et al	2017	US Supreme Court	Meta-factors	28,000	70%
Sulea et al	2017	French Supreme Court	Facts	126,000	96%
Long et al	2018	Chinese civil law courts	Facts, law, applications	100,000	80%
Bull & Steffek	2018	UK Financial Ombudsman Service	Facts	100,000	87%
Chalkidis et al	2019	European Court of Human Rights	Facts	11,500	90%

FIGURE 2.3 AI prediction precision and dispute resolution

where does the case geographically come from? *Second*, what is the legal area of the case? Is it commercial law? Is it criminal law? Is it public law? And so on. *Third*, what is the type of the petitioner? Is the petitioner an employee or is it the State? *Fourth*, what is the type of the respondent? Similarly, is it a corporation, is it an employee, or is it the State? And *fifth*, what is the ideology of the lower court? Is it a liberal court or is it a more conservative lower court? It is important to remember that in this study, the AI did not have access to the facts. It also did not have access to the law. It only had access to meta-factors. But still, the AI showed a precision of 70 per cent in predicting real cases that were decided by the US Supreme Court.

The next study by Sulea et al. concentrated on the French Supreme Court.² They had a bigger dataset. The research team had 126,000 cases, and the dataset in this research project contained facts, but not law. Even though the AI only had the facts of the case, it achieved a very high precision of 96 per cent; i.e. in 96 per cent of cases, the AI was correct in predicting the decisions by the French Supreme Court.

Next, moving on to a Chinese study looking at Chinese civil law courts, Long et al. published a paper in 2018.³ They again had a pretty big dataset of around 100,000 cases. In this project, the dataset contained not only facts but also information on the law and applications. ‘Applications’ refers to the petition, i.e. what the petitioner wanted the court to decide. The AI in this project achieved a precision of 80 per cent.

Furthermore, Bull and Steffek’s paper reported on the Case Crunch Lawyer Challenge.⁴ This study was conducted by Ludwig Bull, who will speak tomorrow, and Felix Steffek, which is me. In this paper, as I mentioned earlier, we received a dataset from the UK Financial Ombudsman Service. Our dataset contained around 100,000 cases, and only the facts of these cases were included. The AI did not have

² Octavia-Maria Sulea and others, ‘Predicting the Law Area and Decisions of French Supreme Court Cases’ (4 August 2017) arXiv:1708.01681v1.

³ Shangbang Long and others, ‘Automatic Judgment Prediction via Legal Reading Comprehension’ (18 September 2018) arXiv:1809.06537v1.

⁴ Ludwig Bull and Felix Steffek, ‘Die Entschlüsselung rechtlicher Konflikte: Der Einsatz künstlicher Intelligenz zur Ermittlung von Entscheidungsfaktoren der Konfliktlösung’ [2018] (5) Zeitschrift für Konfliktmanagement (ZKM) 165.

access to the law. The precision, as I explained earlier when reporting on the Case Crunch Lawyer Challenge, was 87 per cent.

Finally, Chalkidis et al. published a paper in 2019 focusing on the European Court of Human Rights (ECHR).⁵ They had a smaller dataset of 11,500 cases, again containing only the facts. They achieved a prediction precision of 90 per cent.

What can we learn from these studies? First, it is interesting to see the variations in precision. We also get an idea of the number of cases contained in the datasets currently used for such studies. Further, it is quite impressive to see that AI has achieved high prediction precision in many studies, even though it did not get any information on the law. This is quite surprising because, as lawyers, would we not think that the law is essential in predicting the outcome of a case? I am very happy to discuss these issues later if you are interested.

IS THE USE OF AI IN DISPUTE RESOLUTION A GOOD THING?

Steffek: Now I would like to involve you again by asking a question. I would be very grateful again for your answers. My question takes the form of a statement, and I would like to know whether you agree or disagree with the statement. The statement is as follows:

It would be good if citizens and businesses had access to AI services predicting the outcome of dispute resolution procedures such as court litigation, arbitration and ombud schemes.

What is your opinion?

Student D: Access to AI services predicting the outcome of disputes will have a positive impact, of course, but there is a concern that the number of lawsuits may drop a little. I have learned recently that Japan has a low litigation rate compared to other countries, and that is a problem. One of the reasons for this is that the Japanese legal system tends to make it rather easy to predict the outcome of litigation. People tend to give up on their cases before they even start to sue, thinking that their case will probably not be successful. If the number of lawsuits is reduced, it may have the negative effect of reducing the use of the court system and isolating people from legal participation. At the same time, I think that it would be very innovative and useful if it meant that people could anticipate the result of legal procedures, at least as a reference. I would say that I agree to a certain extent, rather than completely agreeing with the proposition.

Student E: I have the same opinion and agree to a certain extent. The reason why I do not fully agree is that I am concerned that if the results are shown in advance by AI, it may cause claims as to why the actual results are different from them. On the other hand, from the citizen's and company's point of view, I believe it has both positive and negative aspects. It certainly makes it easier to plan future actions when damages occur. They will be able to foresee to some extent how much they will have to pay.

⁵ Ilias Chalkidis, Ion Androutsopoulos and Nikolaos Aletras, 'Neural Legal Judgment Prediction in English' (5 June 2019) arXiv:1906.02059v1.

Steffek: Thank you – again, you have identified important questions. Applications of AI – as the studies I have presented show us – are good, but they are not perfect. None of the projects achieves a prediction accuracy of 100 per cent. Hence, the question arises: ‘What if the AI is wrong?’ Your answers also show that you have started thinking about how the world would change if we already knew the outcome of the dispute resolution processes. Would we still go to court? If we knew the outcome in court, then perhaps we would come to an amicable solution to the dispute even before that. Knowing the outcome of the dispute resolution processes would be quite a fundamental change to our societies, our businesses and our relationships. This would also mean that we knew the law, i.e. the legal consequences of our actions. Would it not be great if we all knew the law and the legal consequences that our actions will have? Is there any value in people not knowing the law? These are interesting questions that have been raised by you, and the advances in applying AI to the law make these questions practically relevant.

THE ROLE OF AI IN LAW

Steffek: Figure 2.4 presents the BBC website and how it reported the Case Crunch Competition.

The BBC raised the same issues that you have raised so far. The heading of the website says, ‘The robot lawyers are here – and they’re winning’. In other words, the journalist tells the audience, ‘Look, here are robots that are more capable of predicting case outcomes than lawyers. What does this mean for the human element in justice? Which role should humans play?’ Before we discuss the human element, I will introduce you to further use cases of AI in law. We first need to know how AI can be used in legal services before we can fully discuss ethical issues.

BBC

More than accuracy, costs
and time:

- Further use cases
- Bias
- Transparency
- Ethics




FIGURE 2.4 The robot lawyers are here – and they’re winning

So, how can AI be used in practice? One important point I want to make is that AI can be used on different sides. As shown at the top in Figure 2.5, decision-makers can use AI, i.e. judges, arbitrators and ombudspersons. On the other hand, as shown at the bottom, parties and their advisors can also use AI. The implications are quite different depending on whether a judge uses AI or whether a party uses AI. This is something we need to distinguish when we discuss the desirability of AI applications in the legal domain.

What can AI be used for? It can be used to improve three areas, namely information, analysis and decisions. I would like to give you examples of such use cases. Some of these examples are quite pedestrian and probably uncontroversial, while others are advanced and highly controversial. The fact that using AI in law is sometimes not controversial but sometimes very controversial reminds us that the ethical considerations will depend on the specific use case.

How could a court use AI to improve information? A judge could use AI to find relevant cases in the case database. The judge might be interested in determining what other courts have decided in similar cases, and AI could be used to improve the database search. Or, to give another example, courts could use AI to extract relevant evidence from submitted documents in cases where there are thousands of pages of submitted documents. Again, AI could be used to improve the information available to the court.

How could a court improve its analysis? Courts could analyse past decisions and try to find out whether it has been biased in the past, e.g. by overemphasising certain characteristics of the claimants. How do we do that? Conceptually, it is quite simple. You set up an AI to predict the outcome based on past decisions, and then, you take out one type of information, e.g. the identity of the claimant, and see whether this influences the prediction. This will help you understand whether the identity of the claimant will affect the results. You make a first prediction where the AI has information on the claimant, e.g. age, gender, nationality and whether it is a corporation or a consumer. Then, you ask for a second prediction, but you take this information away from the AI, i.e. you remove it from the dataset. If the




FIGURE 2.5 AI use cases for courts and parties

outcome is different, then it seems that the identity of the claimant may be relevant. Of course, the law applies equally to all, and the identity of the petitioner should generally not be relevant. The law should be administered the same way. If such an analysis of past cases, however, reveals that the identity of the petitioner is relevant in past judgments, it means that there may have been bias in the past. Courts could learn from it and improve current and future decisions.

Third, decisions: for example, you could ask an AI to double-check decisions or get support in drafting decisions. I would like to emphasise that when I describe what AI might do, I am not commenting on whether this is a good or a bad idea. I am only telling you about the possibilities. I say what is feasible. I am not saying that we should do it or should not do it. That is another topic we need to discuss.

How can parties use AI? Parties can use it to identify relevant case law. Advisors can find cases that are similar to the one they are dealing with. In addition, to improve their analysis, they can predict the outcome of going to court or another dispute resolution institution. They can predict the costs of taking a case to court. They can predict how long a case will take in court. As regards decisions, parties or their advisors can ask the AI to help them decide whether to bring a claim to court or to sell the claim. As shown in Figure 2.5, you see ‘access to justice’ as a theme. I think that many use cases of AI can contribute to making law more accessible. This concerns, in particular, the perspective of the parties.

EVALUATING AI FROM A BUSINESS MANAGER’S PERSPECTIVE

Steffek: I would like to hear your opinion again. How would you decide in the following scenario?

If you were the director of a large company, would you use AI for the conflict management of your company? Would you think that it is helpful for you to use it or would you not use it?

Student F: If I were the manager of a large company, I would want to use AI for conflict management. If you are a large company, you are likely to find that devoting human resources to a single lawsuit often interferes with your core business. For example, if you are managing an insurance company, you often have to call an external lawyer for litigation and pay them to attend to the case. By employing AI, you can manage the dispute more efficiently. I believe this would allow us to solve the problem without the reputational harm created by litigation. I think it would be useful to use AI for two reasons: efficiency and reputation.

Student G: I would like to defend the position of not introducing AI. The disputes that large companies have to deal with are likely more complex than, for example, the PPI case mentioned earlier. In the current social climate, there is a risk that society will criticise companies for using AI. In other words, until AI is used in various places and citizens are fully accepting AI, it may be difficult for society to

tolerate the use of AI by large companies for such dispute resolution. As regards complex cases, I would not use AI since the possibility of mistakes that humans would not make cannot be ruled out.

Steffek: Thank you for your excellent contributions. This is a practical question, and you are raising important practical matters. The first comment points to the efficiency advantages that using AI might have. If automating certain legal tasks by way of AI is cheaper than paying someone to do it, then businesses are rational to use AI. This assumes, of course, that the results the AI offers are of the same or better quality. The same logic applies to use cases where AI supports human decision-makers. If AI saves costs or produces better results, then it may be a commercially wise decision to employ it.

The second comment, however, reminds us that AI – at this point in time – might not be able to contribute to all types of tasks. This comment points to very complex cases and doubts that AI is ready to meaningfully contribute to such cases. This is a correct and important point. Currently, the applications of AI in law are quite specific, and there is no AI application in existence that could replace human legal experts in every case. In summary, it depends on the specific task that is to be performed, whether an AI could replace humans in the performance of such a task and, if so, at what cost.

HOW TO SOLVE THE BIAS PROBLEM

Steffek: As mentioned earlier, I would now like to address bias. Bias is a major issue with AI in law, and there are a number of sources for bias. It may be that the dataset has gaps, or there may even be mistakes in the dataset. Perhaps past court decisions were wrong. In such a case, AI will learn from those wrong decisions. Or the algorithm has mistakes and produces wrong predictions. Maybe it chooses the wrong factors, labels the wrong text or is based on a faulty logic. If an algorithm predicts the court outcome correctly in 80 per cent of the cases, it is wrong in 20 per cent of the cases. Or the humans ask the wrong questions to the AI.

What can we do to avoid such mistakes? An important point that I want to make here is that there are technical solutions to bias and there are legal solutions. As lawyers, we may think that law needs to come in and solve the problem of bias, but experts in AI and NLP are also working on solutions. Technical solutions might be to get more data or to de-bias data. De-biasing data concerns, in particular, cases where there is bias in past decisions. Researchers are working on algorithms that help remove bias from data.⁶ Legal solutions, on the other hand, may involve the

⁶ See for example Manish Raghavan and others, 'Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices' (21 June 2019) arXiv:1906.09208; Chalkidis, Androutsopoulos, and Aletras (n 5); critically, Hila Gonen and Yoav Goldberg, 'Lipstick on a Pig: Debiasing Methods Cover Up Systematic Gender Biases in Word Embeddings but Do Not Remove Them' (9 March 2019) arXiv:1903.03862.

- Sources of bias
 - Dataset: gaps, mistakes, bias
 - Algorithm: labelling, coding mistakes, choice of outcome
- Solutions to bias
 - Technical: dataset (more data, less biased data, de-bias data), algorithm
 - Legal: rights and obligations
 - Negative impact v optimal outcome
- Bias in existing dispute resolution systems
 - AI can detect bias in existing dispute procedures
 - Evidence evaluation, decisions
 - AI can provide consistency checks



FIGURE 2.6 Bias

establishment of rights and duties to remove bias. The point that is important here is that there are technical solutions and there are legal solutions. The choice we need to make in specific use cases is whether it is better to apply a technical solution, a legal solution or both.

Looking at the bigger picture, we also need to ask what we would like to achieve as regards the content of AI predictions. Are we interested in optimal outcomes, or are we interested in avoiding negative impact? Often, lawyers focus on avoiding negative impact, but ultimately, I think, AI is a tool to get optimal predictions. Achieving optimal outcomes and avoiding negative impact are not always the same thing, even though an optimal prediction, of course, avoids unwarranted negative impacts. My last point about Figure 2.6 is that existing systems are also not perfect. But AI can help us and courts, for example, to find mistakes or biases in past decisions. This perspective raises the issue of what we compare AI applications against: current practice (which is not always perfect) or optimal outcomes?

THE REMAINING CHALLENGES TO TRANSPARENCY

Steffek: Moving on to transparency brings us to the black box problem that I introduced earlier. It is a technical problem – to start with. When using a neural network, we do not even know what factors the AI uses to achieve its prediction. It is also an ethical problem, since justice is not just about the decision but also about the explanation of the decision. We know this from the way law is taught at university. In the classroom, when discussing legal cases, we do not only ask what the result is and who wins, but we also ask why and discuss the reasons for the outcome. This explanation is part of justice. The same applies to court judgments. They do not only contain the result but also contain the reasons for the result. This is another real-world example making it tangible that justice is not only about the outcome but also about explaining the outcome to those affected by the result.

If AI only tells us the result, that is not the same as justice as there is no explanation. Again, there are technical solutions to the black box problem. Researchers and businesses are working on explainable AI, i.e. they are developing applications that not only predict the outcome but also provide an explanation as to why this outcome is likely. For example, there are research papers using AI to generate explanations.⁷ Such explanations can take the form of the reasons for a judgment or the form of a list of reasons, for example, the applicable statutes and existing cases a court might refer to because they are similar to the case at hand.

As with bias, there are also legal solutions, which can be to require transparency of the algorithm. Another legal approach might be that the parties agree on an algorithm. There are already court cases where parties have agreed on which algorithm to use, and that was a good, consensual solution to the black box problem. Similar to the solutions to bias, the question is whether we use AI to mimic what humans would do or whether we use AI to provide explanations that are better than the explanations humans can provide. To give you an example, let us assume that the average number of reasons for a decision in a first instance court usually provided in a judgment handed to the parties is five. In other words, the parties to a dispute receive five different explanations for the decision of court in the text that is handed to them. Let us also assume that courts provide a qualitative (but not quantitative) explanation of which reasons are more important than the others. Now assume that an AI would be able to generate at least ten or even twenty reasons for the decision. Assume, in addition, that the AI could rank all of these reasons and attribute a percentage figure giving their relative importance in the case. Which type of explanation would you prefer? The qualitative explanation based on five reasons provided by court or the quantitative explanation based on ten or twenty reasons provided by the AI?

USING AI IN COURT

Steffek: This brings me to my last question taking the form of a statement again. The statement is as follows:

As a party I would welcome if the judge deciding my case uses AI to check the judgment (after it is written) and considers amending the judgment if the AI diverges.

In other words, would you like the judge to first decide and then look at the AI and, if the AI decides differently, the judge to check whether they have made a mistake? Would you prefer this to be court practice, or would you prefer the traditional, AI-free approach?

⁷ See, for example Federico Ruggeri and others, 'Detecting and Explaining Unfairness in Consumer Contracts through Memory Networks' (2022) 30 *Artificial Intelligence and Law* 59.

Student H: I am against using AI in such a court setting. I think judgments should be written by humans due to the human and societal aspects of justice.

Student I: I would prefer that AI is used to check judgments and that judges would reconsider their decisions if the AI came to a different conclusion. This is similar to DNA testing, and if there is a technical and scientific basis, then AI should be used to check judgments.

Steffek: Thank you very much for your opinions. I am glad that you voiced different opinions as they show the challenge we face: ‘How much human involvement is needed for a court decision to be just?’ First of all, your answers confirm that different people want different things. The court system and its regulation need to respect this. In this specific case, the solution could, perhaps, be a rule in the procedural code saying that the use of AI to check decisions already taken is only allowed if both parties consent. The question I asked also addresses the point of how good – in the sense of correctness – humans and algorithms are in making legal decisions. Research shows that both humans and algorithms can suffer from bias when determining legal outcomes. It is also worth noting that there is a lot of development in the quality of legal applications of AI. We will also see that the way in which people consider the use of AI will change over time. This tells us that there is no right answer for all people at all times, but that our answer will depend on a number of elements, such as the quality of AI and human decisions, the way in which such decisions are presented (e.g. whether they are properly explained) and individual perceptions of justice.

REGULATING AI IN DISPUTE RESOLUTION

Steffek: I will conclude with a few thoughts on ethics. AI can improve access to justice. However, despite the advantages provided by AI, there are risks as well. Using AI in dispute resolution is not straightforward, and there is no one-size-fits-all approach. I do think that AI will lead to us using law more to create legal relationships *ex ante* and less to litigate *ex post*. This is a good thing because people should know about the consequences of law before they act and not be surprised by the legal consequences later. In the future, we will see more people using technology to understand their legal position before they take action, instead of not understanding legal consequences, then being exposed to a problem later on, going to court and getting a surprise decision. AI will allow businesses and consumers to have a better idea of the law before they make decisions.

Ultimately, it is the interests of citizens that should guide lawmakers, and I think your answers today have shown that there are different citizens who want different things. This is quite a challenge for lawmakers. However, to solve this challenge, we can use existing experience and knowledge. This concerns procedural law, for example. It is not new to us that resources are constrained and that decisions are not always optimal. It is not new that we must find a balance between finding the

truth in court and having constraints in terms of time and resources. Human rights will guide us as regards using AI in law. Procedural rights and principles, too, can guide us. Consumer law can also help us navigate this new field, dealing with situations where one party might have access to AI information and the other might not.

Nonetheless, I think that AI will not just be plugged into the existing justice system, but it will fundamentally transform our legal systems, including legal dispute resolution. I would like to end by saying that in a research project between the University of Cambridge and Hitotsubashi University, we are looking at developing a 'Guide for Regulating Artificial Intelligence in Dispute Resolution'. In this project, we will try to give answers to the questions raised today. I am now at the end of my presentation. It was a great pleasure to have a dialogue with you, and I am very grateful for your opinions. I would like to hand back to Mihoko, and I am very interested to hear what the commentators and further discussion will bring.

WHY JAPAN HAS NOT STEPPED UP ITS USE OF AI

Sumida: Thank you. Professor Felix Steffek leads the team at the University of Cambridge investigating AI in dispute resolution. Professor Kazuhiko Yamamoto is the leader of the mirror working package at Hitotsubashi University, and he is a maestro in civil procedure law. I would like to ask Professor Yamamoto to comment on Professor Steffek's presentation.

Yamamoto: Thank you very much for your very valuable talk. I completely agree with Professor Steffek's concept of utilising AI in dispute resolution. As you mentioned earlier, if parties can access the dispute resolution results predicted by AI, they will be able to resolve their own disputes by, for example, settling them voluntarily, which will result in lower dispute resolution costs and will be beneficial for the parties. It would also be good for society as a whole, as it would allow courts, a valuable resource, to focus on more difficult disputes or more important cases.

However, the preconditions for the use of AI in justice probably differ from one country to another, and what is done in one country will not immediately become possible in other countries. As I see it, there are some differences between the United Kingdom, other Western countries and Japan. I would like to focus on two points.

The *first* issue is the number of court decisions and out-of-court disputes, including those that are known as alternative dispute resolution (ADR), that can serve as the basis of a dataset. Figure 2.3 shows a list of research projects conducted in Europe, the United States and other countries, and the datasets range from 10,000 to over 100,000 cases. In Japan, however, there are generally very few cases of dispute resolution in the public sphere, and more often than not, these cases are resolved in the form of settlements, which are not necessarily public.

Suppose, for example, that you wanted to conduct an AI analysis of a medical lawsuit. In Japan, there are about 800 lawsuits per year for accidents and errors that occur in the medical field. Furthermore, medical lawsuits have a relatively high rate of settlement, with only 200–300 cases resulting in a judgment that is publicly available. If one were to create a dataset as available in other countries – spanning between 10,000 and 100,000 cases –, it would take more than thirty years to collect 10,000 cases and more than 300 years to collect 100,000 cases. That is the current situation in Japan.

The Financial Ombudsman, researched by Professor Steffek, was also established in Japan about ten years ago, modelled after the British example. In Japan, the Financial Ombudsman functions through an organisation called the Japanese Bankers Association, which handles around 100–200 cases per year. The basic policy of financial ADR is to try to resolve disputes with banks through discussion, rather than through court proceedings. As a result, about half of these cases, between 50 and 100, are settled. Therefore, it would take more than 1,000 years to collect 100,000 cases as in Professor Steffek's study.

In addition, in the case of ADR, the Japanese Financial Ombudsman generally does not disclose the details of the cases, as the parties involved very strongly request for the cases to remain confidential. These cases are not accessible to the public. The same applies to many other ADR institutions, e.g. arbitration institutions. Against this background, my first point is that there are limitations in Japan in terms of creating legal datasets.

CIVIL LAW VERSUS CASE LAW JURISDICTIONS

Yamamoto: The second problem is the nature of the legal rules. In Anglo-American law, say, common law countries such as the United Kingdom, where the law is originally based on case law, rules are analysed from the point of view of what exactly is different in each case. This is probably the type of rule system that is familiar to AI analysis.

In contrast, Japan is a country based on a civil law system and codified law. In other words, the system is such that when there are certain legal requirements, certain legal effects will follow clear articles of law. The predictability of the outcome of dispute resolution here is mainly based on such rules. There has been a handful of research in Japan on what kind of legal requirements lead to what kind of legal effects and how to code them into AI such that if certain facts are present, certain legal effects will be produced. In other words, I think we are moving towards coding legal rules rather than analysing actual cases.

Of course, in Japan, too, there are no clear rules regarding, for example, the calculation of the amount of damages. Thus, in the case of traffic accidents, research that analyses precedents and cases as to the damages awarded has been conducted. I understand that this has been done, but I have the impression that this has been the

exception in a rule-based system such as the Japanese. These are some of the differences between Japan and the United Kingdom. However, I believe that the situation in Japan is changing and that it is changing rapidly.

Regarding the first issue of the number of cases in datasets, there is currently a project underway, led by the Ministry of Justice, to make all judgments available in the form of open data. If this is achieved, there is a possibility that at least such judgments could become the basis of datasets for AI analysis. Even though there is the problem of the small number of judgments themselves, I think we will still be able to collect quite a lot of data.

Alternatively, in the ongoing project to digitalise civil litigation, it is being discussed whether all court case records should be converted into electronic data and whether in some cases online access should be allowed. If this can be achieved, there is a possibility to utilise not only the judgments but also the litigation records that form the foundation for the judgments as a dataset.

However, there are certain hurdles to these projects in Japan. Regarding the conversion of judgments into open data, there is a strong need to anonymise the parties involved. Currently, AI is used to remove names from case records. However, many parties in Japan may hesitate to go to court if all related case data will be available on the internet.

Similarly, criticism has begun to emerge against the idea of allowing internet access to the litigation records mentioned earlier, as it may result in an infringement of the right to a fair trial. From this perspective, it seems that in Japan, public perception is different in some respects from that of other countries in the West. I believe that this issue is closely related to developing datasets for conflict resolution.

THE CHALLENGE OF UNLOCKING THE BLACK BOX

Yamamoto: Regarding the second issue of a civil-law-based system, I think that rapid changes are taking place. With changing social conditions, rules need to be as flexible as possible, and as I mentioned earlier, there is a move towards developing rules that can respond to changes in society by stipulating general clauses, using abstract concepts. In such a legal system, it is important to understand what factors judges specifically consider when applying such general rules, clauses and requirements, how they weigh the relevant factors, which factors they consider important and which factors have less relevance. When applying more general rules, there needs to be clarity on which factors are important and which are not so important. AI can contribute to this.

It is not easy to manually analyse the weighting of different factors in such cases, but I believe that the 'hidden layer' described in Professor Steffek's Figure 2.2 can provide new insights for analysing such rules. This is also related to the issue of transparency in AI, which Professor Steffek mentioned at the end of his talk. If the

weight of such factors can be analysed, it will be of great significance to the parties concerned, to the judges who will make decisions based on them and to the legislators who are responsible for developing the rules. Hence, I have great expectations for this project, and I hope that the students who will be responsible for the future will look at jurisprudence from a new perspective. The relevant discussions will be very meaningful.

Sumida: Thank you very much. Today, we also have a member of the Ministry of Justice with us, Counsellor Fujita. Counsellor Fujita, I would like to ask you to contribute your thoughts.

CAN JAPANESE LAWYERS ACCEPT AI?

Fujita: My name is Fujita and I am from the Civil Affairs Bureau of the Ministry of Justice. First of all, Professor Steffek, thank you very much for your very interesting talk on AI. I would now like to make a few comments, including from the perspective of a civil judge, although I myself am currently working for the Ministry of Justice and have been involved in IT in the judicial field. I am a practitioner, so I am in a position to lose my job as AI advances, which is a prospect I am not very keen on (laughs). Despite that, I believe it is a meaningful project that meets the needs of the times, and I would like to give a practitioner's opinion.

I would like to talk about three perspectives that I feel should be considered in relation to dispute resolution when new technologies, including AI, emerge, in the light of the current situation in Japan. The first concerns the technological potential of AI; the second investigates the mindset of the lawyers who are in a position to use it; and the third focuses on the public's acceptance of the services enhanced by technology and the security of the users.

Regarding the first point, the technological potential, in recent years, the potential and usefulness of using new technologies, including AI, has steadily increased in the legal sector, such as for contract verification and legal research. Many researchers and practitioners have become interested in this field. As a recent trend, the implementation of various service offerings, including those offered by start-up companies, is also progressing in Japan. There, apart from infrastructural aspects such as the small amount of data available in Japan, restrictions and difficulties in terms of technology are not much heard of. In the legal services sector, it can be seen that more advanced use of technology is possible, at least from a technical point of view. In the case of Japan, the challenges faced by the customer side, the lawyers and users, may instead be greater.

On the second point, concerning the mindset of lawyers as regards AI, I was joking earlier when I said that as a lawyer, I do not want AI to take my job. I personally think that human lawyers and legal AI can coexist well with an appropriate division of roles. Currently, in Japan, I have the impression that many lawyers want to use IT and AI as tools to assist and streamline their work but are

afraid that it will replace them in their role as a lawyer who thinks and analyses. I think this fear is unfounded.

Why do Japanese lawyers have such an attitude? There are two possible causes – first, the traditional working style of Japanese lawyers. Japanese lawyers often go through several stages of listening to the client's story, analysing it, relating it to legal texts and precedents, drawing conclusions and seeking solutions while managing the client's expectations. The process is often carried out in several stages, moving back and forth between the lawyer and the client. It can be said to be a process that is tailored to the individual client and the dispute. Such a process, even if not consciously arranged, is likely to be a highly individualised approach. This creates the impression that such a service cannot be replaced by something automated or mechanical.

The second cause is rooted in the nature of the legal profession. Japanese lawyers may hesitate to prioritise convenience when it comes to the core of their legal services. This is because convenience often entails some risk. In Japan, there are very strict institutional rules and restrictions on the provision of legal services by persons who are not qualified as legal professionals. This practice has long been entrenched in the system. It is a tradition to which lawyers, judges and other legal professionals have been exposed. They all resolve legal disputes with high ethical standards and personal responsibility. They take pride in providing legal services themselves as skilled experts, without relying on external technology, such as AI.

Against this background, discussions concerning the use of IT and AI in the justice sector often raise the question of who is liable in the event of a technical fault or unexpected malfunction. Even if AI makes an obvious mistake due to unintended information processing, most lawyers are probably comfortable with the position that lawyers are not exempt from the risks and liabilities associated with the use of AI. The risk aversion of Japanese lawyers may be one of the reasons why the use of AI in the legal system has not progressed.

WHAT DO USERS WANT FROM DISPUTE RESOLUTION?

Fujita: The third perspective concerns us as the users of the legal system. It is necessary to rethink how the parties involved in a dispute can feel satisfied with the resolution of the dispute. Japan has traditionally been said to have a strong 'sense of superiority', i.e. the public authorities are considered great, and they are thought to protect the common people. In an emergency, people want to rely on absolute authority. In legal disputes, a certain expectation and trust seems to have been placed on specialists such as lawyers or on having disputes resolved through the open, official authority of courts. Even if it takes some time and effort, such an open and clean resolution is considered to be a good thing.

However, recent changes in social conditions and technological developments have changed the situation of disputants. Looking at e-commerce and contracts on

the internet between individuals, we wonder whether our traditional idea of judicial resolution of disputes through experts is always the right thing to do or whether it is the best way for all parties to gain a sense of satisfaction. Many people might rather prefer a more convenient and quicker resolution. As Professor Steffek said, it may be that many people will now prefer a resolution service that can be obtained easily and free of charge on the internet, rather than having to spend money and time consulting a lawyer. We have to rethink what good dispute resolution means for the parties concerned. From our point of view as practitioners, it may be necessary to change the mindset a little and to take an approach to dispute resolution that is easier and excludes emotional conflicts as well as the individuality of the case to a certain extent. Also, we should consider the possibility of closed solutions, such as private ADR, rather than public, rights-based solutions, such as court litigation. We should also think about a more open and transparent approach to dispute resolution, rather than a face-to-face solution. Furthermore, I wonder what kind of effect chat-based rather than face-to-face interactions will have in the future. There, seemingly contradictory interests, such as convenience and due process, will have to be reconciled and balanced.

One final point – as Professor Yamamoto mentioned earlier, the Ministry of Justice and others have in recent years begun to make concrete moves towards the digitalisation of court procedures and the application of open data principles to judgments. Given the time constraints today, I hope to talk about this at length if an opportunity arises later. To conclude, in Japan, judicial innovations are still on their way, in the process of being developed, and in this context, I listened to today's discussion with great interest. Thank you very much.

Sumida: Thank you very much for your valuable comments. Professor Steffek, would you like to respond to the Japanese side's comments?

CITIZENS ARE AT THE CENTRE OF JUSTICE

Steffek: Thank you very much, Professor Yamamoto and Counsellor Fujita. These are very interesting and thoughtful points. I would like to concentrate my response to give us more time for discussion and thus would first like to comment on Professor Yamamoto's ideas that I found very fascinating. Professor Yamamoto raised the issue of the number of cases needed in the training dataset. I think that is, indeed, a very important point. Currently, AI research is also looking at whether good results can be produced with a lower number of cases in the dataset. The existing research papers tend to use datasets with many cases because such datasets are more likely to achieve better results. But I think we will see in the future whether, for example, a few hundred cases that cover a certain area can also produce good results. Research is not concluded as regards this issue. Listening to Professor Yamamoto, this will be quite important for jurisdictions such as Japan where there are not so many cases available in the legal system.

I would also like to comment on Professor Yamamoto's point where he said that in civil law systems, the law plays a bigger role, whereas in case law systems, as in Britain, court judgments are relatively more important. That is true. In the research project carried out with Hitotsubashi University, we will have an opportunity to compare the application of AI to Japanese law and English law. I expect that we can gain important insights from this comparative research, and I am very grateful that we have this research collaboration between the two institutions.

Thank you very much, Counsellor Fujita, for your very interesting thoughts and comments. I very much agree with the point you made that most practitioners expect that AI will supplement but not replace them. It seems to me that the situation in the United Kingdom and other countries is similar. Concentrating on tasks such as database research, providing advice, developing innovative contractual solutions, etc., is more helpful than only asking whether technology will replace human lawyers or not. The more granular focus on tasks is more precise. I also think that we can learn from past technology and how it has transformed the legal profession. How have computers changed the legal profession? What impact have email and the internet had on the legal profession? Technology has always transformed the legal profession, but it has never replaced lawyers. Computers and email have not replaced lawyers, but they have changed the way in which lawyers provide their services. That is why I think that this conversation series is a great opportunity for us as participants, since, as Professor Yamamoto said, we can learn and adapt to new ways in which law is applied. I think that, say, in ten or twenty years, legal services will look fundamentally different from what we see today. I also think your point, Counsellor Fujita, of helping people to develop trust in technological advances is very important. We need to look at the people and what they want. It is for the people to say, and that is the beauty of justice. It is about the people and the citizens and their wishes. But of course, not all wishes can come true, and this is the challenge we face.

Sumida: Thank you very much. Professor Takeshita is a member of the project who is also tuning in today. Professor Takeshita, would you like to add something?

CRITICALLY EXAMINE YOUR OWN COMMON SENSE

Takeshita: My name is Takeshita and I am based at Hitotsubashi University. Thank you very much, Professor Steffek, for a really excellent lecture. I specialise in private international law, and I am interested in AI in relation to the use of online dispute resolution (ODR) in international dispute resolution. In the light of Professor Steffek's lecture today, I would like to comment on the significance of the use of AI in jurisprudence.

Dispute resolution prediction by AI has considerably overturned the conventional method of decision-making applied in Japanese civil jurisprudence. Although there are various types of AI, the AI for dispute resolution prediction introduced by

Professor Steffek inputs only facts and predicts the resolution of disputes as an output. Usually, if you make legal judgments in Japanese civil jurisprudence, as Professor Yamamoto mentioned briefly, you apply legal norms that combine legal requirements, *Tatbestand*, and legal effects. You determine in a step-by-step manner whether the legal requirements have been met, and then you conclude whether the claim is finally upheld or not. It is a very interesting point to skip all those legal judgments and let the AI learn and come to a conclusion in a way that links the facts and the conclusion.

It is interesting that in one study reported by Professor Steffek, AI has only learned from meta-factors, which is not an approach used in the normal application of the law. Furthermore, based on the results of the Chinese project, it appears that better results are achieved when only the facts are input rather than the AI having access to both facts and the law. I also thought that this was very interesting.

I would like the students to think about the following: if an AI that learns only facts and conclusions can predict dispute resolution more accurately than a human being who carefully applies the law, what is the application of the law then good for? Personally, I think that the significance of using AI's dispute resolution prediction ability is to critically examine the ideas of jurisprudence that we consider common sense. We can then see whether the application of law that we are undertaking is really the correct way of going about it.

There can be a real diversity of ideas in jurisprudence, and AI-based dispute resolution prediction will force people to reconsider the ideas in jurisprudence that you think are common sense. Nevertheless, as Professor Steffek pointed out, AI decisions still have a black box problem, i.e. that it is not directly revealed as to why a given solution is reached.

It is indeed a black box situation, where AI predicts a conclusion of conflict resolution with a high probability of being 'right', but the users do not know what the reasons or underlying factors are. In relation to this point, what I would like students to bear in mind is that if a certain meaning is to be given to AI prediction, including whether it is 'correct' or not, it still has to be done by a human being. For example, if the parties to a dispute demand an explanation for the outcome of the dispute resolution, there may of course be the possibility of an AI that automatically creates an explanation, as Professor Steffek mentioned. For the time being, however, a human hand will still be required to provide an explanation. In other words, the process of interpreting and applying the law that we are engaged in could be said to be exactly such a process of explanation and understanding, whereby it is assessed whether the human interpretation and application of the law is appropriate, when confronted with the results of dispute resolution presented by AI. Critical examination must be carried out, including the predicted results of AI learning from human-initiated cases, and it seems to me that it is still only humans who can do this.

Furthermore, AI basically learns from cases that have been judged in the past and predicts dispute resolution. It, thus, still may need human hands to realise a developed interpretation of the law. I believe that the significance of research in this field for the time being is to make efforts to ensure that humans properly understand the meaning of dispute resolution predicted by AI and to utilise this in human justice, as Professor Sumida also talked about the ‘coexistence of AI and humans’ in one of her previous classes.

Sumida: Thank you very much. Felix, would you like to say anything after listening to all those comments?

CO-EXISTENCE WITH AI REFINES HUMAN INTELLIGENCE

Steffek: Yes – thank you very much, Professor Takeshita. You mentioned that the research papers I have presented show that approaches where the AI had access to only the facts produced results with high precision. I think we could spend a whole session discussing all these papers and why they achieved a higher or lower precision. Such precision also depends – often quite substantially – on data pre-processing. In some of the existing studies, in particular, studies that I have not presented today, the question arises whether the facts already contained signals from the judges about what the outcome would be. Hence, the AI may actually have found clues in the facts about what the outcome would be. This would not be desirable because, ideally, we would want facts in the datasets that do not indicate the outcome.

To conclude, on the one hand, we realise that AI is not perfect. There are no AI applications in law with a prediction accuracy of 100 per cent. On the other hand, the question is whether the existing justice institutions, e.g. courts, are always 100 per cent correct. Some cases go to appeal and are overturned, so it seems that not all decisions by courts are perfect. The question then is how we can use AI to make justice better and avoid the risks that AI brings. Thank you very much, Professor Takeshita, for your insightful comments.

Sumida: Would anyone like to speak? Professor Takeshita also raised some issues with the students, and if anyone was inspired in any way by today’s discussion, I would be very grateful if you could tell us if you have changed your mind.

Student J: Yes, thank you very much for today. I am mainly studying statutes on a daily basis, but I was reminded of the importance of case law. It seems that the process of drawing conclusions in AI is based on case law, so I will put more effort into the same. It is just a brief impression, but thank you very much.

Sumida: In the previous session, I was talking about an analogy, i.e. that the coexistence of humans and AI has been realised in the world of shogi, Japanese chess, for a while now, and that what is happening in the world of shogi will, to some extent, also happen in the world of law. In the shogi world, amazing tactics are developed by AI through a process completely different from that of humans. The

royal road of tactics, in which it had been said that ‘whoever masters these tactics will win the shogi world’, has lost against new moves developed by software. The world of shogi and go in Japan is a leading example of humans improving their intelligence while coexisting with AI. Japan is at the stage where research such as that introduced by Professor Steffek today is still in its infancy, but shogi is a world of professionals that is familiar to Japanese people.

DIFFICULTIES IN PROCESSING JAPANESE LEGAL DOCUMENTS WITH AI

Yamada: Artificial intelligence covers quite a wide field of research. Famous applications are in the areas of robotics and gaming, including chess and shogi. Also, fundamental studies such as search, inference and machine learning are important. Among AI research topics, NLP is an essential field in AI and law. Natural language refers to the language we humans use. NLP is a field of research to develop technologies that allow computers to process human language. NLP aims to save human effort by, for example, analysing large amounts of documents efficiently by computational tools or generating sentences by computers. Typical applications include *kana-kanji* conversion (Japanese input method), machine translation and document retrieval. In the context of Japanese legal NLP, the judgment retrieval system has been the most widely used application. However, now that the scope of application of NLP in the legal field is expanding, as we see, the ‘legal tech’ concept has become more and more popular in recent years. The application of NLP to the legal field has become a field in which many research institutes and companies are working worldwide.

CHALLENGES IN PROCESSING LEGAL DOCUMENTS WRITTEN IN JAPANESE

Yamada: In NLP research, the target of studies is text data. As English and Japanese are different languages, NLP systems developed for each language are also different. Most of the NLP systems are developed in English, and there are fewer Japanese studies than English ones. When it comes to legal NLP, this tendency is more remarkable. Is it difficult to create a system for the Japanese language, especially in the legal domain? Actually, we can utilise NLP techniques developed in other languages (often in English) even for the Japanese legal domain.

Most of the methods proposed in recent research are data-driven, i.e. they involve learning patterns based on the collected data, rather than rule-based methods, which process according to predefined rules. In data-driven methods, input sentences are divided into words (or similar units such as morphemes) and converted into computer-processable representations by assigning a numerical representation to each unit. Thus, in subsequent stages of learning and reasoning, the differences

between languages are not a critical barrier in transferring techniques from one language to another. Even if a method is initially developed for English, it can be used for Japanese as long as enough Japanese data are available for learning. We cannot guarantee that every system developed for English can easily be adapted to Japanese without any issues. Still, the principal methods can often be transferred and utilised for different languages. Therefore, the unique characteristics of the Japanese language are not critical obstacles.

What about the uniqueness of the legal domain? Does it hinder the application of NLP? There are differences between linguistic expressions used in everyday life and those used in legal documents. For example, the vocabulary used in news documents naturally differs from that used in contracts. These differences can be resolved in a similar manner as the differences between languages if enough data in the legal domain are available for learning.

However, there might be technical issues caused by the characteristics of the Japanese legal system and legal documents, which are more complex than simple differences in vocabulary and expressions. There are many areas where research and development must be conducted. Nevertheless, some of the issues can be solved by applying existing NLP technology where sufficient data are available.

LEGAL NLP AND DATA AVAILABILITY: OUTSIDE JAPAN

Yamada: Due to data availability, legal NLP research is more popular abroad than in Japan. In the legal domain, judgment documents are more readily available in electronic and machine-readable formats in the United States and the United Kingdom, whose justice systems are based on the common law. They are far ahead of the Japanese legal system. The machine-readable format refers not just to an electronic version of a document but to a format of data authored in a structure that allows a computer to extract and process its contents easily.

The Caselaw Access Project (CAP) in the United States has published a database containing all precedents in the United States published up to around June 2018, which is not simply scanned data in PDF format but data available in a machine-readable format.⁸ In Europe, case data are also easily accessible through the EUR-Lex⁹ and ECHR websites.¹⁰

These publicly available datasets help apply NLP technology to the field of law. For example, a language model called BERT,¹¹ which has been widely used since

⁸ Caselaw Access Project (Harvard Law School) <<https://case.law/>> accessed 1 November 2023.

⁹ <<https://eur-lex.europa.eu/>> accessed 1 November 2023.

¹⁰ HUDOC: European Court of Human Rights <<https://hudoc.echr.coe.int/>> accessed 1 November 2023.

¹¹ Jacob Devlin and others, 'BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding' (11 October 2018) arXiv:1810.04805.

2019 for its performance and versatility, was adapted to the legal domain in 2020,¹² using data from CAP and ECHR. Better data availability enables rapid development and application of NLP in the legal domain. Data availability is also crucial in terms of experiments and evaluation, as well as the reproducibility of research.

APPLICATIONS OF NLP AND DATA AVAILABILITY IN JAPAN

Yamada: Various applications of legal NLP have already been implemented in Japan. As a support tool for legislation and administration, the Japanese Ministry of Internal Affairs and Communications provides the e-Legislative Activity and Work Support System (e-LAWS), studied by Nagoya University's Toyama Laboratory.¹³ e-LAWS provides consistent support from the drafting and deliberation of bills to their publication on the web. The system stores legal data in a machine-readable and unified format. Thus, it is also a source of open data in the Japanese legal sector. In the field of corporate law, NLP has often been used in contract analysis and drafting. Corporate law work, which is based on text-based records such as contracts or agreements, would be a good target for NLP applications. Recent advances in NLP technology and the rise of needs in corporate digital transformation have made legal NLP applications more popular.

On the other hand, in Japan, NLP applications for documents related to courts (e.g. judgments) appear to be less advanced than abroad. There is a wide range of possible applications for legal NLP in this area, e.g. in improving judgment retrieval systems, pseudonymisation of court decisions and automated summarisation. However, we need the documents to be publicly available in a machine-readable form. Currently, only a small number of judgment documents are available on Japanese court websites, and a machine-readable format has not been employed. Compared to the previously mentioned examples, such as the CAP project and the ECHR, it is much harder to obtain judgment document data in Japan.

If the data are unavailable, it is difficult to conduct even basic experiments using the existing technology before developing advanced systems and methods to reflect the characteristics of the Japanese legal system and documents. In order to discuss the benefits, challenges and concerns that new technologies such as NLP and AI will bring to the judiciary, it is necessary to evaluate their possibilities through actual research and development. To this end, the first step is improving the availability of the data, ideally as open data. In Japan, discussions are underway to make civil judgments open data.¹⁴ This should be achieved as soon as possible.

¹² Ilias Chalkidis, 'LEGAL-BERT: The Muppets Straight Out of Law School' (6 October 2020) arXiv:2010.02559.

¹³ <www.kl.itc.nagoya-u.ac.jp/index_e.html> accessed 1 November 2023.

¹⁴ <www.jlf.or.jp/wp-content/uploads/2021/04/pt-houkoku20210325.pdf> (in Japanese) accessed 1 November 2023.

CONCLUDING CONVERSATION

Professor Felix Steffek and Professor Mihoko Sumida

Sumida: Professor Steffek, the man behind this series of conversations, appeared on stage to start the discussion. The discussion got off to a great start with the students actively speaking up. The issue of ‘technology changing conflict resolution’ is at the heart of our joint research project between the University of Cambridge and Hitotsubashi University. Thank you very much for sharing your insights.

Steffek: It was a great experience to exchange views with the members of the research project and the Japanese Ministry of Justice, and, of course, to interact with the Japanese students. I have never taken a class like this before, but if I were a student, I would definitely enjoy it (laughs).

Sumida: I totally agree with you on that (laughs). I was pleasantly surprised by the engaged responses of the students to my bold attempt to turn a challenging piece of research into an educational programme for the next generation. I thought that Figure 2.5 was very innovative. It clearly shows the world view that the use of AI in dispute resolution can qualitatively improve the information, analysis and judgment of both ‘fair and neutral decision-makers’ and ‘disputants’, and this will lead to the realisation of ‘access to justice for all’ for the society as a whole.

Steffek: Certainly, the message of the session was that we wanted people to understand that there are myriad ways in which AI can contribute to improving access to justice, and I certainly consider Figure 2.5 very important. It shows that the first contribution of AI to dispute resolution concerns ‘information’: AI can improve the quality of the collection and retrieval of legal information by, for example, locating relevant and significant legislation and case law. Secondly, AI can improve the quality of the ‘analysis’: it can assist with the examination of available information and provide the factors that may lead to a favourable or a non-favourable outcome in dispute resolution. Thirdly, AI can assist in ‘decision-making’: not only can AI be used in the decision-making process, but it can also assist in the drafting of judgments. Public and academic debate often revolves around this third area of innovation specifically because it is both the most glamorous and morally challenging. The contribution of AI to the acquisition and analysis of information may be modest, but it is very important for the practice of dispute resolution.

I also wanted to show that there is a use case for AI not only for neutral decision-makers such as judges, arbitrators and ombudspersons but also for parties to undergo a more effective and efficient dispute resolution. Thus far, it is private parties, often commercially based actors, who have used AI to manage disputes. The slow adoption of AI by neutral decision-makers is not surprising given that the use of data and algorithms by authorities is far less accepted than in the private sector and raises difficult ethical questions.

Sumida: The research project currently involves AI researchers to develop a predictive tool for resolving civil disputes, and in the second half of the research project, you plan to develop ethical guidelines for the use of AI in dispute resolution.

Steffek: It is important that lawyers are also familiar with technology and discuss it together with AI researchers. Some of the well-known risks and challenges of using AI include the risk of bias in predicting results and the issue of transparency in understanding the factors taken into account by AI – the ‘black box problem’. The solutions to these problems lie in both technology – improvements in data and algorithms – and law – the imposition of rights and obligations. Lawyers in particular need to be aware that not all risks require legal solutions. For example, research is underway on de-biasing data and on explainable AI technologies that provide better transparency.

AI has great potential to improve access to justice, but its risks need to be understood and managed. The final and deciding factors are the ‘interests of the citizen’, which must be kept at the centre of measures undertaken by legislators, while taking into account human rights, procedural rights and research on the balance between the benefits and challenges of AI.

Sumida: I would like to focus on the word ‘justice’ in access to justice for a moment. In Japan, ‘access to justice’ seems to have been associated with a rather specific idea, such as support for law enforcement overseas, or even within Japan, legal advice in areas affected by large-scale natural disasters. However, if you think about it, the idea of using technology to spread the understanding of law in civil society and to ‘realise justice through law’ is not unique or special. It should be a normal issue for us, and I feel that the world has realised this again after the pandemic.

Steffek: I was happy to see that both Professor Kazuhiko Yamamoto and Counsellor Fujita of the Ministry of Justice were very welcoming to the idea of improving access to justice by introducing AI. As Counsellor Fujita stressed, the key is to ensure the public’s trust in the judiciary. This may include dealing with cases where the AI has made incorrect predictions. For example, we may need to consider a new appeals mechanism.

Sumida: The last point is a very interesting one. In Japan, there is a consumer group litigation system that allows unspecified consumers to sue for their interests. Should we consider an AI screening version of such a system?¹⁵ In the Netherlands, a human rights group sued to stop the use of an AI algorithm to detect welfare fraud by pointing out its bias.¹⁶ Nevertheless, I am very curious about who identified the

¹⁵ This direction is proposed by Souichirou Kozuka, *Governance of AI Development and Use* (in Japanese), NBL 1150, p. 31 (2019).

¹⁶ Jenny Gesley, ‘Netherlands: Court Prohibits Government’s Use of AI Software to Detect Welfare Fraud’ (*Law Library of Congress*, 13 March 2020) <www.loc.gov/item/global-legal-monitor/2020-03-13/netherlands-court-prohibits-governments-use-of-ai-software-to-detect-welfare-fraud/> accessed 1 November 2023.

problem with AI algorithms in the Netherlands and how they established the issue. And why were human rights organisations able to request an injunction against the algorithms? This is another area that needs to be researched.

As data scientist Cathy O’Neil says, it is very difficult to identify AI problems and challenge them in court.¹⁷ I also mentioned consumer group litigation because in a Japanese society where the internet has become so widespread, most people are now considered to be ‘consumers’ of online platform services. There seems to be an opportunity to take consumer protection as a starting point. On the other hand, the insights from the Netherlands story cannot be ignored. If we are to connect the two, there is a need to radically and significantly upgrade the concept of consumer protection. First of all, the concept of the interests of consumers who are harmed may need to be reconsidered. For instance, it might be problematic if there are data scientists among those supporting collective litigation. There are likely to be many other issues, which we will continue to consider.

QUESTIONS FOR FURTHER THOUGHT

- What would a world look like, in which everyone has access to dispute resolution outcome predictions? How would dispute behaviour and management change? What would be the effect on commerce or on consumers? Would this be a better or a worse world?
- What are the consequences if not everyone has access to dispute resolution outcome predictions? Do parties today have the same information available before going to court?
- Should all court decisions be publicly available for everyone at no cost? Should certain elements of the judgments (such as names) be deleted or anonymised before making them available?
- What is special as regards decision-makers, i.e. those deciding on the outcome of dispute resolution such as judges, arbitrators and ombudspersons, using AI? Does the administration of justice require a human decision-maker? What is the difference between AI being used before and after the human decision-maker comes to a conclusion?
- Does the use of AI by decision-makers hinder the development of the law as AI learns only from past judgments?

¹⁷ Cathy O’Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (Kubo Naoko tr, Crown 2016).