

Irène Buvat and Ken Herrmann Talk with Alexander Stremitzer, Kevin Tobia, and Aileen Nielsen

Irène Buvat, PhD, CNRS director of research and head of the Inserm Laboratory of Translational Imaging in Oncology at Institut Curie, Orsay (France), and Ken Herrmann, MD, MBA, a professor of nuclear medicine at the Universitätsklinikum Essen (Germany), talked with Alexander Stremitzer, Kevin Tobia, and Aileen Nielsen, who recently published a paper on the liability of physicians who use AI systems in the *Journal of Nuclear Medicine*.

Buvat You have recently published in our journal a very scientific paper on a quite unusual topic for our journal. Could you please introduce yourself?

Stremitzer I am a lawyer and economist. I used to work at UCLA until two years ago and I am now at ETH Zürich running an interdisciplinary group of lawyers and social scientists, psychologists, philosophers, even including a physicist working on legal matters and running experiments on legal institutions.

Tobia I am an assistant professor at Georgetown. My background is in law and philosophy, as well as experimental cognitive science. I am also an affiliate and former member of Alex's group at the Center for Law and Economics at ETHZ.

Nielsen I am Aileen Nielsen, a member of Alex's group, ETH Zurich. I have done previous studies in physics and anthropology, as well as being a U.S. trained lawyer. I run experimental studies looking at the interplay of law and technology.

Buvat Could you please summarize your study and its main findings?

Stremitzer Sure. As you know, often AI precision tools make personalized treatment recommendations, such as recommending the dosage of the drug based on the patient's file. This is the example that we use in the paper. Medical malpractice requires a deviation from the reasonable care standard and the reasonable care standard is traditionally met if the physician exercised standard care. This raises the possibility that the physician exposes himself to liability when following nonstandard advice by the AI. This would, of course, undermine the very promise of precision medicine, which is to offer personalized and therefore nonstandard advice. This is a theoretical possibility, but we also saw a countervailing force, which was that maybe accepting the AI advice is increasingly viewed as the new standard. How the legal system weighs these points is largely an empirical question. Because it is a new thing, we cannot really observe lots of court decisions. Therefore, we decided to exploit the fact that such decisions, at least in the U.S., are decided by jurors who are laypeople. So we randomly assigned a representative sample of U.S. adults to different scenarios to study how laypeople judge the reasonableness in different scenarios. In those scenarios, the AI either recommended standard or nonstandard care, and the physician either accepted or rejected that advice but in all scenarios, it turned out that the patient was wrongly treated. We found that laypeople want doctors to basically follow both the standard treatment and the AI advice. In short, we find that physicians that accept standard advice are exposed to considerably less liability if something goes wrong than a physician who rejects standard advice. But the really interesting point is that a physician having accepted nonstandard advice or rejecting nonstandard advice is almost equally likely to be liable if something goes wrong with a slight advantage for accepting the nonstandard advice. This suggests that laypeople acting as jurors are far more open to AI recommendations than one would have thought.

However, if a physician receives nonstandard advice, he is less protected than if you receive standard advice and accepts it. So rejecting standard AI advice seems to be a very bad idea.

Buvat Was that what you expected?

Tobia We went to the study with an open mind with the idea of attempting to adjudicate among several different models of how laypeople would make these kind of liability judgments. Yet, a prior article in JAMA by Price et al had conducted a doctrinal analysis and that analysis formed part of our background expectations, which was that the current state of the law favors that physicians follow standard care. But that article also suggests that this might change quickly. So perhaps in the future, physicians will incur liability for rejecting correct AI nonstandard recommendations. So, I guess one way of thinking about our results is that they suggest that we are actually a bit further down that line than the scholars expect. We are sort of already in the future. One final reason that there is something surprising about that finding is that there is a lot of other research that found that laypeople are averse to algorithms in certain contexts. Here we found actually that when those two factors conflict, when the AI recommendation is for nonstandard care, actually accepting that advice was evaluated as more reasonable than rejecting it, which is in some sense surprising.

Herrmann My conclusion is follow AI advice. If AI advice is standard, I am lucky because I am not going to be sued anyway. And if I am unlucky and it is not standard, I have a higher likelihood to be sued. But still less than if I reject it. Correct?

Stremitzer We basically find that this the recommendation at the current state of the tort system and it goes against the intuition of this doctrinal analysis in the JAMA piece.

Buvat The population that you surveyed was from the United States. Do you think the results might be different in Europe or in Asia, because the citizens might not have the same view on artificial intelligence algorithms?

Stremitzer Well, that is certainly possible. We tried to get a representative sample of U.S. adults, because in the U.S., the legal system relies on the judgment of laypeople. One major difference in Europe is that the legal system relies on professional judges. I think that might even be a bigger difference than intercultural differences. So, we could easily run the study on laypeople in different countries, but I think what would even be more interesting is run a similar study on some European judges. These are, of course, much more difficult subjects to recruit. It could be that there is a difference in how laypeople, medical professionals, and the legal elite view those things.

Buvat Do you think that when the physician does not agree with what the AI suggests, this should be mentioned in the medical report?

Stremitzer It might be that the AI advice and the doctor's judgment differ. I think it would be hugely valuable if doctors documented that even if they decide to follow AI advice because they feel it has some advantages for liability. They might also be humble and say who am I that I know better in that situation? If everybody systematically documented differences, you would actually be able to see areas where the AI advice and the doctors' judgments deviate. I think exploring the sources of these divergences would probably be very interesting for AI precision tool development.

Buvat Might the liability of the physician be different if he gave some explanations associated with his final decision to follow or not the AI recommendations?

Tobia I think it is likely to either be neutral or probably help the physicians' case in terms of lay judgment. That is something that would be an interesting further question to investigate with a similar empirical study to see whether laypeople are responsive in their liability judgments to the fact that more justification or explanation has been offered.

Buvat Can't the AI algorithms just be considered as a different type of medical experts? If so, what does that change regarding the liability of the physicians?

Stremitzer If the AI is taking more and more responsibility away from the doctor, it would be conceivable that in the future there is going to be liability of tool developers as there is product liability for other products. So the liability might be shared or even, the responsibility shifted from the doctor to the physician to the tool. But this is kind of pure speculation.

Tobia We have another study that addresses this question a little bit. We take up this really nice strategy that you suggested of thinking about certain AI essentially as a colleague. Making this kind of comparison to see whether how people evaluate an AI as someone the physician might rely on versus a human that a physician might rely on. The project is still on-going, I want to leave it as a teaser. Maybe over time people will start to see AI as more like a colleague and less like a tool. Building on prior work, we found that one thing that predicts people's reasonable judgments is their conception of what is common. If AI is seen as more common in practice, people will increasingly likely hold physicians not liable for relying on AI advice.

Herrmann The acceptance of AI is, based on your study, already quite groundbreaking. If you talk to people who provide these systems, they only say they provide you with a decision support system. But based on your results, I would say in laymen's terms, it is not even accepted as a decision equal, it is actually accepted as a superior, because if you do not follow an AI, you are screwed ! For me, the doctor going to school, having fifteen years of experience, this is actually completely counter-intuitive !

Stremitzer There is still the judgment of the doctor because there is another factor for being held liable. If there is no harm, you will never be held liable. So, if you have a very strong reason to believe that the AI is telling you the wrong thing, you can actually act against that on the theory that you will just influence the probability of harm. And yes, there is a slight advantage of following the AI, even if you get nonstandard advice. But I think probably it is not as strong as when you look at the standard advice.

Buvat Your discussion about standard care is extremely interesting, because we are in the era of precision medicine. Precision medicine means that each patient is unique and that you have to customize the best treatment for each patient, so the notion of standard care might become less present. Do you think the physician liability will thus get increased because there might be less and less standard care?

Nielsen Whether precision medicine increases the risk of liability will depend on the population and how do those particular populations perceive personalization. It is true that our experiment suggests that receiving standard advice and accepting that advice creates a relatively strong shield against liability. This might not be this clear as treatments become more and more specialized. But as Alex mentioned before, it is not clear to what extent the responsibility will stay with the doctor. So, in tandem with this drive towards

personalization, is increasing reliance potentially on AI products, in which case we can expect more liability to also be allocated in the form of product liability to app developers and scientists who are developing the baseline algorithms.

Buvat So this is good for physicians. They will not get sued any longer!

Nielsen Right. So that would be having your cake and eating it, too, that you have less liability. But still, full authority and decision-making power.

Buvat I was wondering whether you think physicians should be trained more on AI to have a better understanding the ways these things work and make the best use of them. Or as long as they are FDA or CE approved, the physicians will rely on them anyway ?

Nielsen What we as lawyers can tell you is thinking about how FDA approval works and how that should affect physicians. So, what we know is that FDA approval establishes a bare legal minimum for sale in a medical market. It actually does not in any way shield you from the possibility of tort liability for medical malpractice or for product liability. Clearly, physicians necessarily need to know more when they are considering liability implications, just as with a medication or a medical test or a medical device. Mere knowledge of FDA approval is neither enough to operate responsibly nor enough to make liability decisions. But the extent of that education and the shape of that education very much need to be determined with experimental and empirical studies

Herrmann But shouldn't the layman be trained or understand how the algorithms work because they seem to trust them quite a lot.

Tobia That is a great question. We are studying laypeople in part as potential patients, but largely as potential jurors who might be making a decision about what was reasonable. In a real trial, there is also going to be other evidence beyond what we have in our study. There will be expert testimony about the device that is used. So, this also highlights one limitation. So, I think there are a number of good reasons to have that education of laypeople in terms of how medical AI is and could be used. As a baseline, it would help in certain contexts, securing more meaningful informed consent or coordinating expectations between the patients and physicians. It is an open question whether a physician or medical experts view of AI is similar to what we found in the lay sample or different. Empirical work and also education can help coordinate those sets of expectations.

Buvat As lawyers, do you see any danger in the spread of AI algorithms in the healthcare system?

Nielsen Absolutely. Of course, it is our job as lawyers to see potential problems or dangers. Rubber stamping is a concern. We want to make sure that we keep human physicians engaged, especially as they will continue to have meaningful exposure to legal liability. We want to make sure that their decision-making authority continues to be proportionate to that. Thinking about the data science element, we also know that existing data driven methods do not necessarily have all the strengths that human physicians have. We know that there is a danger of runaway feedback loops. We know there is a danger of bias data. So we very much see a place for humans in the loop to avoid these dangers. People might think that these problems are solved. But I would point out just at the end of 2019 in Nature, there was a publication of a major U.S. insurance company rolling out an unintentionally biased algorithm, a racially biased algorithm, because they use medical spending as a proxy for seriousness of medical condition. Of course, unfortunately,

medical spending has a great deal of racial bias built into it. So even a very sophisticated national health insurance company unintentionally rolled out bias. It is very important to recognize that this continues to be a problem and having humans in the loop may be helpful. We can also wonder about the balance of innovation versus safety. Laypeople are perhaps so gung ho but not educated enough about these. We can worry that there would be some sort of public relations push, for example, by hospitals or competitive pressures to adapt these tools perhaps before they are fully vetted or appropriate. Likewise, regulators are just getting into the business of regulating AI products. The FDA, for example, said more than 10 years ago that they would regulate these products or had the authority to do so. But there have not been that many products through the process. So, unlike any medication or devices where they have decades and decades of experience, this is clearly going to be a new learning experience, even for very sophisticated regulators to think about how to balance innovation with patient safety. And finally, we want to think about a balanced perspective and encouraging that balanced perspective in many stakeholder populations. So, we want the layperson patients, the physicians, hospital administrators, regulators and app developers to be skeptical, but not too skeptical. So, these are the dangers and how we can try to avoid them as we promote the proper regulation and legal regime.

Buvat Can you tell us about the advantages that you see from the lawyer point of view?

Nielsen Many of the advantages from a legal perspective are in fact solving some of the dangers. Bias in medicine is certainly not new due to AI. We know that our physicians are unfortunately and unintentionally sometimes producing unintentional bias in the system. You can see AI as an opportunity to reduce that. Just as humans can complement machine learning weaknesses, machines can complement human weaknesses. Humans are not necessarily good at high dimensional data. Humans are not necessarily good when they get tired. And these are the sorts of things that can lead to legal liability. When a human is looking at a file and may be overwhelmed because there are so many inputs or human is simply working overtime, so may not physiologically be in the best position to make decisions, they will nonetheless face legal liability. They now have a machine that will perhaps help them prevent some of these errors. So there are certainly opportunities for growth as far as the legal perspective on medical malpractice. But ultimately, I would say we are a little bit out of our depth because we would also be looking to scientists and physicians to tell us where the opportunities for improvement are on the ground. Then we can think about how to make sure the legal system sets the appropriate incentives for adopting those when they mature and get appropriate for mainstream use.

Buvat and Herrmann Thank you all for your time, your input in that passionate “AI for healthcare” debate is extremely valuable.