

EXPERT REMARKS
CCDH EXCHANGE OF VIEWS – DRAFTING GROUP ON THE USE OF NEW
TECHNOLOGIES BY THE EUROPEAN COURT OF HUMAN RIGHTS

VERONIKA FIKFAK, 11 May 2026

I have been working on the issue of automation in human rights adjudication for the past six years. Most of what I will present is based on an article published in the Michigan Journal of International Law with Laurence Helfer entitled, *Automating Human Rights Adjudication* and my work since then.¹

In my talk, I would like to make four points, four recommendations. The first is to underline how different human rights adjudication is to normal domestic adjudication. The second is to emphasise that the suitability of automation or AI for the Court's processes may depend on the type of automation and artificial intelligence tools that are being contemplated. Put simply, certain questions can and certain cannot be entrusted to the machine. The third point concerns how we approach the introduction of such tools into the adjudication process and what we can expect of those using these tools – the “human-in-the-loop”. Humans have often been considered as a check on the machine. But the reality is that humans are biased and increasingly over reliant, even dependent on AI. We thus have to put checks and balances in place to ensure that the interaction of human and machine does not amplify or introduce new biases. My final point will turn to the issue of accountability and transparency. What responsibilities does the Court have towards the applicants who turn to it and how transparent should it be about its use of AI and automation tools?

On my first point, I would like to highlight the specific characteristics of European human rights adjudication, which make it different to domestic adjudication. They thus require of us to be even more careful when introducing automation or AI into the adjudicative process.

When individuals turn to international courts and tribunals, they do this only after they have exhausted all domestic remedies available. This means that they turn to the international court as a matter of last resort. In addition, when the European Court of Human Rights renders its judgment, these decisions are final and individuals cannot challenge them. Put together, this means that if the Court with the assistance of automation tools reaches decisions that are flawed but that remain final, it can mistakenly deny individuals their right to receive appropriate remedy.

¹ Veronika Fikfak & Laurence R. Helfer, *Automating International Human Rights Adjudication*, 46 Mich. J. Int'l L. 69 (2025). Available at: <https://repository.law.umich.edu/mjil/vol46/iss1/3>.

The second characteristic I should underline is that European Court of Human Rights adopts an expansive approach to rights interpretation – human rights law is constantly evolving as it is applied to new situations like climate change. This dynamic, evolutive approach directly challenges some of the inherent elements of automation and AI – namely its static nature. Because of this status quo bias, such technologies could prevent or slow down the evolutive interpretation of the Convention.

We thus have to think very carefully about how these technologies are used at international level. This is even more the case, since decisions rendered by the Court have an influence across 46 countries of the Council of Europe. All of these elements – the last resort, the finality, the living tree doctrine, and the cross border application have to inform our approach to what use of new technologies before the ECtHR we allow.

My second point relates to the suitability of automation and AI use in human rights adjudication. The court has adopted a number of measures to make its work more efficient. These can include scanning applications or allowing them to be sent by email (from certain jurisdictions); digitalising the text of the applications to make them searchable and flagging up missing information; using summarisation and translation software; calculating damages in excel spreadsheets; content and similarity clustering of repetitive cases to ensure they are handled in a similar manner; using decision-tree question prompts on issues like admissibility and repetitive issues. Most of this automation is done for the purposes of digitization and for case management to facilitate and speed up processes. These uses are generally uncontroversial.²

The greatest concern we should have, however, is about the introduction of AI tools to facilitate judicial decision-making. This can include AI to facilitate research, problem solving, or judicial reasoning. For example, there is currently a plan to create a new AI program Chatbot– a sort of a ChatGPT for ECHR case law – which would help Registry lawyers and judges answer (legal) questions from the Court’s case law.

Looking at the *letter* and *spirit* of ECHR, which stipulate that a decision has to be rendered by an independent and impartial tribunal, it would be highly

² Although there are certain concerns about the inaccuracy of translation tools and their incapacity to appropriately capture emotions and affect. Satterthwaite (2025) notes that LLMs are often poorly trained on small – so-called low resource – languages, which raises risks when AI is used for seemingly simple tasks. Indeed, given that most litigation is fact-intensive, in that facts are often more determinative of the outcome than a legal interpretation, these risks should not be underemphasised. Satterthwaite, M. (2025). *Report of the Special Rapporteur on the independence of judges and lawyers: Artificial intelligence in judicial systems: promises and pitfalls*: UN doc. A/80/169, 16 July 2025.

In addition, translation tools – like Deepl – often dull emotional, passionate language, which is inherent in human rights complaints and there is therefore concern about the extent to which it can capture the complaint made.

problematic to substitute the discretionary decisions of judges with even the most sophisticated AI. We should reject the use of AI to predict violations or decide admissibility issues involving high degrees of discretion, such as whether a claim is “sufficiently substantiated” or “manifestly ill-founded.” But we should also be careful about use of AI for research purposes or indeed to aid with judicial reasoning.

There are several aspects, for example, in thinking about building Court’s own ChatGPT, that we may wish to consider. The aim of such a system is for lawyers and judges to ask legal questions to which the ChatBot can provide an answer and suggest relevant jurisprudence. Such an inventory system could raise several problems: first, these ChatBots require large datasets to learn from. We may think that we have a rich and expansive dataset from which our ChatBot can learn, but when broken down into specific violations and issues, the ECtHR case law is not that big. If we think only about slavery or forced labour, it turns out we only have a handful of cases to draw on. Second, similar ChatGPT models used by LexisNexis and Thomson Reuters hallucinate between 17 and 33% of time. Such hallucinations are more frequent especially when not enough case law is available. Any suggestion provided by ChatBot would thus need to be doublechecked against the case law. The third issue is how the results are presented by ChatBot: in response to a question, the ChatBot may provide a formal presentation in support of a particular outcome, yet it may hide the particular facts of the case that explain the choice of the decision. Since these are crucial for the understanding of the outcome, using the output presented by the tool may lead to its incorrect application.

This then brings me to the third point and that relates to how humans interact with AI, the so-called human-in-the loop question.

Traditionally, we think of human-in-the-loop as a corrector of a flawed algorithm: an AI generates a recommendation, if it gets it wrong, a human will fix it. But studies have shown that this is not the case and that humans can hardly act as supervisors of AI recommendations. In practice, humans suffer from automation bias, whereby they *either* blindly follow the algorithm, rubber stamping its recommendations without considering how they were reached, or they automatically discount its recommendations when in fact these should be followed. The presence of this bias is directly tied to the trust they have in the technology as well as to who is ultimately accountable for the final decision. Individuals may see AI tools as a moral buffer, a distancing device which diminishes their responsibility. ECtHR judges and lawyers thus have to stay accountable for their decisions.

Beyond this, however, we also want to worry about what impact the use of an AI tool may have in the long term. Excessive use can of course lead to overdependence on AI. Studies have shown for example, university teaching, since the introduction of ChatGPT, students' work has become more mediocre, lacking depth and often engagement with subject matter. The concerns about over-reliance on AI are raised also in the health field because they hinder the development of critical thinking and original writing skills.³ In a judicial setting, these are crucial in the context of decision-making and reasoning. But reliance on AI outputs invites a growing dependence on algorithms that lack genuine comprehension of legal precedents, statutes, and contextual nuances.

Training, disclosures, and cautionary instructions can redress some of these concerns, but only if they focus on providing a better understanding of how recommender AI systems work and what their limits are. When individuals understand how algorithms are integrated into institutional decision-making processes and what they can and cannot do, their knowledge of how to interact with the tool improves, leading to enhanced trust and decreased bias.

On a systemic level, therefore, before ChatBot or similar tools are introduced, experiments should be carried out to evaluate human-algorithm interactions. These experiments should investigate questions such as: "Who is the specific human engaging with the algorithm and how are they doing this? Are there gaps in knowledge that could limit their ability to assess outputs and anticipate concerns (eg automation bias, overreliance)? To what extent might the algorithm curb human discretion or assessment that is essential to the decision (eg not providing all the facts)? Who are the other human actors responsible for shaping the system?" The answers to these questions will help to determine how the tool should be incorporated into the decision-making processes.

This then leads me directly to my final point on transparency, explainability and accountability.

How should we ensure the Court remains accountable to its users? More specifically, how transparent should we be about the information we provide about the AI tools that are used in the process and what explanation should we provide to the applicants in individual cases?

³ Zhang S, Zhao X, Zhou T, Kim JH. Do you have AI dependency? The roles of academic self-efficacy, academic stress, and performance expectations on problematic AI usage behavior. *Int J Educational Technol High Educ*. 2024;21(1). <https://doi.org/10.1186/s41239-024-00467-0>. Also in health: Rosas-Jiménez, C.A., Lokker, C., Ibhawoh, B. *et al*. Trust and overreliance on ChatGPT and its implications on critical thinking: an exploration in health policy. *Discov Public Health* **23**, 200 (2026). <https://doi.org/10.1186/s12982-026-01568-z>.

The European Court of Human Rights should follow a *structured process* for evaluating *whether* to introduce automated decision-making and *how* automation tools – including AI – are introduced. The process should begin with a human rights impact assessment that should (1) explain how automation tools are appropriate to the Court’s goals and will be integrated into its decision-making processes, (2) describe the tools’ anticipated benefits, the foreseeable risks involved, and the measures for mitigating them; and (3) indicate how judges and staff will be trained, including what data has been or will be gathered in the testing period to understand the interaction of lawyers with such recommendations and their ultimate impact on judicial outcomes. This is not a productivity assessment, but rather an evaluation of how human-machine interaction might be changing the adjudication process and the success or failure of applications.

The second component of accountability focuses on the *explanations* provided when automation is used. Courts should provide *two types of information*: (1) a *general* explanation of the institutional oversight process that I just described – eg by publishing the human rights assessment and the training undertaken and its impact on the outcomes; and (2) *case-specific* information to *applicants* about whether automation/AI was used in their case and in what way. Taken together, these disclosures and explanations serve to ensure that the process remains fair and that values of individual justice that are central to human rights adjudication are upheld.

It is imperative that there is transparency about what is already in place and what will be introduced. If we move too fast, this will have implications. For example, if *flawed* decisions reached with the assistance of AI become precedents for future cases against the *same country*, against *other states parties*, and even for *other* regional human rights systems, this is likely to generate calls for redress and challenge of decisions that are already final. Indeed, some scholars already now argue in favour of a *right to challenge* a decision made by or with the help of automation or argue for *a right to have a human* make the decision in the first instance (in effect, this is a right to opt out of automation altogether).

Such calls, which are likely to only grow stronger with time, open up a dilemma for human rights adjudication. An important reason for introducing automation is to expedite the processing of cases. A right to challenge decisions reached with the assistance of automation tools, or a right to opt out of those tools, would require creating multiple tracks for reviewing claims, significantly increasing time and resources. It would also impose additional burdens on international courts that are significantly over-stretched and under-resourced. It is thus better to proceed with automation slowly and transparently, whilst testing human-machine interactions and automation’s impact on outcomes.

To conclude therefore:

Automation and AI raise foundational questions about (1) *how* individual rights claims are adjudicated and what in this process is suitable for automation and what not, (2) *how* humans interact with machines, and (3) *how* we can ensure international institutions remain accountable to their users. These questions require of us to proceed slowly and with caution so as to ensure the necessary checks and balances are incorporated into the system and able to redress any errors.

Thank you.