



Oxford Institute  
of Technology and Justice

PERSPECTIVES

RUDIMENTARY AI IN JUDICIAL  
DECISION-MAKING:  
LESSONS FROM COLOMBIA

Alejandro Moya Riveros

Blavatnik School of Government, University of Oxford



**CLOONEY  
FOUNDATION  
FOR JUSTICE**

The views and opinions expressed herein are those of the author and do not necessarily reflect the official position of the Institute.

Copyright for this work remains with the authors.

# RUDIMENTARY AI IN JUDICIAL DECISION-MAKING: LESSONS FROM COLOMBIA

By Alejandro Moya<sup>1</sup>

I am a Colombian lawyer and philosopher, and I worked as an Auxiliary Judge at Colombia's Special Jurisdiction for Peace, known as the JEP by its Spanish acronym. It is a transitional court that combines restorative and retributive justice in cases of war crimes, crimes against humanity, and genocide committed during the country's armed conflict.

In what follows, I will address the "rudimentary" use of AI in the courtroom—something I explored at the JEP. I will start by telling you where I come from (since that background shaped the decisions I later made) and what I did with AI.

The JEP<sup>2</sup> was established in late 2016, when Colombia signed a peace agreement<sup>3</sup> with the FARC, a guerrilla group that had fought for over 50 years. At the time, I was living in the United States, finishing my master's studies at Harvard Law School. Soon after graduation, I was looking for a good job in Boston or New York—something juicy, ideally to help pay my student loans. But the peace agreement changed everything.

That agreement filled many Colombians with hope, especially those of us who had worked on the armed conflict and knew first-hand how crucial it was to end it. The negotiations with the FARC and the creation of institutions like the JEP made us believe a different Colombia was possible. So much so that I dropped my original plan of working abroad and flew straight back home to find a place at the JEP.

Luckily, my old constitutional law professor and dean had just been appointed as a magistrate in the JEP's Appeals Chamber. I contacted him immediately and asked for a job. He appointed me as his legal officer. We began working in April 2018. I was 28.

Six months into the job, I was fascinated by the cases I worked on but couldn't help feeling frustrated. My boss reserved the most sensitive legal discussions for himself and his two

---

<sup>1</sup> LLB & BA, Universidad de los Andes; LLM, Harvard University; MPP, University of Oxford; MSc in Public Policy Research, University of Oxford (candidate). Associate Researcher and Judicial Specialist of Oxford Institute of Technology and Justice. Email: luis.moyariveros@bsg.ox.ac.uk

<sup>2</sup> See more on the JEP at <https://www.jep.gov.co>.

<sup>3</sup> See more on the Final Peace Agreement at <https://peaceaccords.nd.edu/wp-content/uploads/2020/02/Colombian-Peace-Agreement-English-Translation.pdf>

auxiliary judges. I longed to join those conversations—to contribute, as that was the reason I had returned to Colombia in the first place. One day, I gathered enough courage, stumbled into his office, and asked for a promotion. I wanted to be a judge too, though I barely met the legal requirements.

He took a month to decide. Finally, he said yes, perhaps because he saw in me the same ambition he once had. I became the youngest auxiliary judge, both in the Appeals Chamber and in the entire court. All my peers were in their forties or fifties. I was not yet thirty.

Very quickly, I realised how unprepared I was. I lacked both the knowledge and the experience to be a good judge, and everyone could see it. Many times, I felt like an impostor. Deep down, all I wanted was to do my part, but often I didn't simply know how. At first, no one in the Appeals Chamber took me very seriously, even when I raised the right argument. When I circulated draft opinions, they attracted more criticism than anyone else's. I often felt powerless. Sometimes, I regretted asking for the promotion at all.

It took me years of hard work to feel comfortable in the role, to make real contributions, and to earn the respect of my colleagues. I worked over 13 hours a day, with almost no weekends or holidays. I risked, most likely unwisely, my social life, my relationship with Ana—my beloved and understanding partner—, my health, and my peace of mind, almost to the point of burnout.

In November 2022, OpenAI launched ChatGPT-3.5 as the first public chatbot powered by a large language model. The JEP, though innovative in its judicial approach and fully digital after the pandemic, offered no guidance or position on AI use. Some of us in the Appeals Chamber were curious, but using AI in our work was risky. We were under constant scrutiny from the media and the government; some judges and auxiliary judges had even been threatened, attacked, followed, or spied on. Engaging with AI in judicial decision-making could add another layer of danger.

Still, I was intrigued. There remained a gap between me and some of my peers, and AI seemed like a useful tool to work faster and better. More importantly, I shared the growing concern that the JEP was running out of time to complete its short 15-year mandate. To get the job done, I believed, the institution needed every resource available. So, I decided to give it a try.

Very soon, I started noticing the difference. In the second semester of 2023 and the first of 2024, I drafted what I like to believe were three of the five most important decisions of the Appeals Chamber in that period. Usually, a case of the magnitude—and difficulty—of

those I helped resolving, took me about a year to understand and draft. Now I was cracking it in roughly four months. But beyond speed, those rulings stood out for their deeper understanding of the legal system, their stronger awareness and respect for the JEP’s own case law, their creativity and boldness in interpretation, and their clearer structure and language. Far from falling behind my peers, I was leading the way. And the Appeals Chamber—and the JEP more broadly—benefited from those efforts.

I used AI in what today would be called *supportive judicial decision-making*—not to decide cases, nor to have AI suggest particular outcomes, but to assist me in legal research, analysis, and writing. All of this was done with due caution: never sharing personal or sensitive data of the parties involved, directly verifying any information provided by AI (especially regarding norms and case law), and never asking it to write, on its own, any part of a decision. I authored every sentence in the drafts that later became rulings, following independent—and most likely analogue—review by my peers and senior judges.

With the above caveats, AI helped me to:

1. **Be more coherent.** By uploading key, previously published decisions into the chatbot, I was able to build upon them—not merely align with them or avoid contradictions. One of Ronald Dworkin’s most famous metaphors is the “chain novel”:<sup>4</sup> the idea that the law, like a novel, is written in successive chapters, each different, but contributing to a shared story. Instead of restating what the Appeals Chamber had already ruled, I aimed to push its interpretations forward and develop the underlying normative theory.
2. **Achieve deeper, more nuanced understanding.** Large language models don’t have all the answers—least of all in law, which is too specific and complex. But they can serve as sparring partners. Judging is a lonely profession, especially in places like Colombia, where the judge is often a full-time writer. And, as García Márquez once confessed, nobody helps you to write your stuff.<sup>5</sup> There you are, alone, in “absolute loneliness”, facing the case files and the blank page. Although you eventually submit your work to peers or senior judges, collective deliberation only happens at the end, once the first draft is finished. AI became my constant companion, testing my reasoning, refining my arguments, and helping me foresee alternative interpretations and counterarguments.

---

<sup>4</sup> See Dworkin, Ronald. *Law’s Empire*. Cambridge: Belknap, 1986. Print. Pg. 228 ff.

<sup>5</sup> See García Márquez, Gabriel. *El olor de la guayaba*. Bogotá: Debolsillo, 2024. Print. Pg. 118.

3. **Tackle issues beyond the strictly legal.** Judges often need to engage with disciplines and tasks outside the law—like designing or refining procedures that encourage wrongdoers to acknowledge responsibility, devising ways to repair harm, or fostering victim participation. These are challenges a behavioural economist, a psychologist, an anthropologist, or a professional manager or policy maker might handle with greater ease. AI helped me grasp the basics of such problems—where to look, what to read, and how to talk about them.
4. **Develop a consistent institutional voice.** This is crucial in a collective body where parties and the public expect decisions that sound coherent, as if written by a single author. AI also helped me achieve greater clarity—a key element in transitional justice, where the purpose is not merely to resolve cases but to build peace through judicial narratives that everyone can understand.
5. **Save time, and use it better.** Judges, like all professionals, have limited time. They must reach timely decisions, no matter how complex the case. AI allowed me to save energy on repetitive or mechanical tasks and reinvest it in thinking more carefully about the substance of my cases. Better time management led to better case law.

By using AI in this way, I worked faster and achieved better results. But let me be very clear: it demanded more intense intellectual engagement. If there is a moral to this story, it is that AI is not a shortcut to our jobs. It is not just about efficiency or convenience. Used strategically and responsibly, it can improve the quality of justice itself. It can make judicial systems work as they are meant to—and even better. And that is what people deserve, especially in times like this, when the law often feels like society's last resort.

Retributive criminal law stopped doing the trick long ago. Harsh punishments and high conviction rates do not deter crime, nor do they help victims heal. Unsurprisingly, few are satisfied with justice today. But we won't get different results by doing the same things. That, in my view, was the premise of Colombia's Peace Agreement and the JEP: the belief that even the gravest crimes could be met with restorative justice. Using AI, therefore, is not just a right or a privilege for judges—it is, I am tempted to say, a duty.

Those of us who have worked in the judiciary often believe in the law's power to fix problems, no matter how complex. We see ourselves as bureaucrats in the good—perhaps British—sense: guardians of the public interest. And to be frank, those sitting in high courts rarely lose sleep over backlogs or inconsistency. What keeps us awake at night is the fear that the judicial system stops making a real, positive contribution to society, that it becomes

an empty cult of the law, producing paper rulings no one obeys, unable to change lives, or worse, harming people in the very process meant to protect or rehabilitate them.

So, the real question when addressing AI in judicial decision-making is not whether it will make us more efficient or consistent—it will. The question is what new and higher purposes can we now pursue with the power of AI? Which of our judicial aspirations, once impossible, are now within reach?

Yet, if we are to pursue these new purposes, we must first confront AI's limits. Predicting the "the right" judicial decision based on datasets of cases and norms—which is what some private developers are doing today—may be statistically accurate but not necessarily legally correct. The law follows its own logic—and that logic lives in the judge's mind. Some call it "judicial intuition": the deep, tacit knowledge judges accumulate after resolving hundreds of cases, allowing them not merely to detect patterns, but to grasp the spirit, the *telos*, the theory of justice itself. Yet this knowledge is unwritten. It appears in every decision, but never in full form. Perhaps because no single case demands its complete articulation, or because it is hard to express. For that reason, judicial intuition is not part of any dataset. And AI cannot learn from it—at least not yet.

If AI cannot yet grasp the theory of justice that lives in a judge's mind, then perhaps it is time we make that theory explicit. We would all benefit, as AI would then assist judicial decision-making not by replicating patterns, but by feeding from a deeper understanding of the law. That is precisely where judges must step in. Instead of becoming spectators to the evolution of AI—or worse, to our own obsolescence—we must claim our place at the table and bring something new to it. We must assume a new role: to make explicit the theory of justice that guides our work. To abstract further, make the judicial apparatus intelligible, and write it down. And then to refine it continuously. That would be our greatest contribution to justice in the age of AI, and what will keep us relevant, valuable, and profoundly human as technology continues to evolve.