



INTERVIEW WITH PROFESSOR RODRIGO CANALLI ON THE USE OF ARTIFICIAL INTELLIGENCE IN THE JUDICIARY

INTERVIEWED

Rodrigo Canalli is a Law clerk to a Justice of the Superior Labor Court (TST) and Professor at CEUB. Master in Law, State and Constitution from University of Brasília (UnB). Master in Competition, Innovation and Information Law from New York University School of Law. Former law clerk to a Justice of the Supreme Federal Court and Head-Advisor at the Artificial Intelligence Advisory Office within the same court.

INTERVIEWERS

Christine Oliveira Peter da Silva

Doctor and Master of Laws in Law, State and Constitution from the University of Brasília. Professor in the Master's and Doctoral Programs in Law and Public Policy at the University Center of Brasília. Secretary-General of the Constitutional Studies Center of the Brazilian Supreme Federal Court.

Ian Ferrare Meier

READ CEUB's Editor-in-Chief. Professor in the Lato Sensu Postgraduate Program at CEUB. Researcher. Master's student of Law and Public Policies at CEUB.

Luísa Cristina Vasconcelos Marimon Álvares

READ CEUB's Editor-in-Chief. Lawyer. Postgraduate student in New Tendencies in Public Law at CEUB. Undergraduate student in Philosophy at UnB. Law graduate from CEUB.

Lucas de Pádua Carvalho Mendes Lima

Member of READ CEUB. Undergraduate student in Law at CEUB. Has internship experience in the Supreme Federal Court and four law offices.



Interview with Professor Rodrigo Canalli on Artificial Intelligence in the Brazilian Judicial Branch^{1 2}

Interviewers:

Christine Oliveira Peter da Silva³

Ian Ferrare Meier⁴

Luísa Cristina Vasconcelos Marimon Álvares⁵

Lucas de Pádua Carvalho Mendes Lima⁶

How to cite this interview: CANALLI, Rodrigo. Interview with Professor Rodrigo Canalli on Artificial Intelligence in the Brazilian Judicial Branch. Interview granted to Christine Oliveira Peter da Silva, Ian Ferrare Meier, Luísa Cristina Vasconcelos Marimon Álvares and Lucas de Pádua Carvalho Mendes Lima. **Revista de Egressos e Acadêmicos de Direito do CEUB**, Brasília, v. 1, ed. 1, p. 89-101, 2026. Available at: [...]. Accessed on: [...].

USE OF ARTIFICIAL INTELLIGENCE BY THE JUDICIARY

[Interviewers] — How can the relation between artificial intelligence and law be described? Which challenges does law face in the face of the new technologies, and how is AI being applied in judicial practice?

¹ Law clerk to a Justice of the Superior Labor Court (TST) and Professor at CEUB. Master in Law, State and Constitution from University of Brasília (UnB). Master in Competition, Innovation and Information Law from New York University School of Law. Former law clerk to a Justice of the Supreme Federal Court and Head-Advisor at the Artificial Intelligence Advisory Office within the same court.

² The interview herein provided is the product of the participation of Professor Rodrigo Canalli in an event of the Centro Brasileiro de Estudos Constitucionais (CBEC) of the Centro Universitário de Brasília (CEUB).

³ Doctor and Master of Laws in Law, State and Constitution from the University of Brasília. Professor in the Master's and Doctoral Programs in Law and Public Policy at the University Center of Brasília. Secretary-General of the Constitutional Studies Center of the Brazilian Supreme Federal Court.

⁴ READ CEUB's Editor-in-Chief. Professor in the Lato Sensu Postgraduate Program at CEUB. Researcher. Master's student of Law and Public Policies at CEUB.

⁵ READ CEUB's Editor-in-Chief. Lawyer. Postgraduate student in New Tendencies in Public Law at CEUB. Undergraduate student in Philosophy at UnB. Law graduate from CEUB.

⁶ Member of READ CEUB. Undergraduate student in Law at CEUB. Has internship experience in the Supreme Federal Court and four law offices.

Rodrigo Canalli — The connection between artificial intelligence and Law is characterized by a strong interaction. Here, the Law must offer the answers or, at least, the questions related to social, economic, and legal matters that emerge in face of a new technology that enters the scene and is popularized.

A clear example of this is in the Supreme Federal Court, which has held a public hearing to discuss algorithmic subordination, a legal question brought forward by a recent technology that didn't exist. Furthermore, there's a noticeable peak in discussions on copyright, due to generative artificial intelligence tools that produce images, videos, and texts, raising new problems to which copyright law wasn't previously prepared and now must provide answers to. And recently, the use of generative artificial intelligence in election campaigns has also demanded action from the Superior Electoral Court, which had to regulate situations that weren't imagined before, because they didn't exist.

On the other hand, we must also point to the use of artificial intelligence in the legal field. By that I mean its incorporation by lawyers, who use these tools for procedural management, precedents' research and writing of filings. Additionally, the Judicial Branch also uses these technologies to aid in its proceedings, in the management of lawsuits collections, assisted judicial decisions or automation of procedures, amongst other applications.

[Interviewers] — Could you explain a bit more about the inner workings of Artificial Intelligence within the legal field?

Rodrigo Canalli — When addressing the incorporation of artificial intelligence in law, we are talking about a specific type of computer algorithms, that is, instructions destined to be interpreted by a machine, which possesses specific characteristics. What are those characteristics? Generally, when we talk about artificial intelligence, we mean technologies that are capable of recognizing patterns and producing decisions in situations of uncertainty, based on statistical criteria. Additionally, they are enhanced through use, an idea that summarizes the essence of machine learning. Machine learning is a type of algorithm that generates progressively improved models, based on previous experience, through trial and error.

The programming of artificial intelligence, differently from traditional programming, which is based on binary concepts such as “yes or no,” “zero or one,” “all or nothing” — that being

called Boolean logic —, makes extensive use of statistics and mathematical statistics. We now have a more complex logic, as it incorporates statistical reasoning.

Another analogy I like to make is as follows: when addressing computer codes, algorithm means instruction. It is a set of instructions in a language that will be understood by a computer, nothing beyond that in essence. If we consider a cake recipe, it can be seen as an algorithm in a language interpreted by a person to produce the result, which is a cake. Similarly, the same principle is applied to a civil procedures code, or a criminal procedures code. It's a set of instructions written in a language destined to be interpreted by judges and lawyers, who, by following them, will produce a result, which is a correct and legitimate legal decision.

We are talking about code and instructions. There's a natural affinity in employing these types of analytical tools in our professional activity, that being the law. This is one of the ways in which it is most evident.

[Interviewers] — In which area has Artificial Intelligence stood out the most in the Judicial Branch?

Rodrigo Canalli — One of the areas in which artificial intelligence has stood out significantly when it comes to law is on the identification and analysis of precedents. But what does identifying precedents mean? Analyzing precedents is nothing more than recognizing patterns, which demands the ability to abstract a case and generalize a hypothesis. That is basically what we teach to language models through machine learning training, known as natural language processing. This process consists in teaching a machine how to comprehend our language, learn the lexicon, the rules of syntax, as well as the rules of context and semantics, to solve the ambiguities that are inherent to language.

Language models, such as GPT-4, which is behind ChatGPT, Meta's AI or Google's Gemini, are the motors of these tools. They are designed to learn from an extensive language universe. In the field of law, we'll have specific models developed by courts, universities, or private sector partners, increasingly focused on comprehending the legal language and extracting relevant information, including the identification, within a specific case, if it is or isn't covered by a binding precedent.

And why has the Judiciary been so interested in artificial intelligence? With each new technology that emerges, there's a glimpse of an opportunity of attempting to solve the Judiciary's great problem: the huge caseload. A few years ago, when digital procedure was implemented, it was believed that it would end the stock of paper cases. However, the challenge persists with the accumulation of digital archives. Now, artificial intelligence also brings forth those expectations, through the idea that artificial intelligence will help in achieving faster and better decisions. It's possible that it might truly help in that sense. It's worth remembering that according to recent data from the National Council of Justice (CNJ), there are about 80 million cases pending judgment. In this scenario, it's evident that a technological resource with such a potential must be explored with all the necessary precautions, so that we may benefit from it fully, in the best way possible.

[Interviewers] — Which are the main benefits of the use of artificial intelligence in the Judiciary Branch?

Rodrigo Canalli — One of the main benefits within the use of artificial intelligence in the Judiciary is the amplification of the capacity for rational management of the case backlog. This can be observed, for example, in the Supreme Federal Court, thanks to the VitorIA, a tool which assists by managing the case backlog and, consequently, accelerating internal procedures and the court's judicial activity. The Superior Court of Justice also has, in a larger scale, a tool that performs the same function of amplifying the capabilities for rational case backlog management, besides other courts that make use of diverse similar tools.

Another significant benefit of artificial intelligence for the Judiciary Branch is its ability to contribute to the production of more consistent sentences, allowing the application of the old adage of "treat like cases alike." Although the pattern identification might, eventually, contain errors, it has the potential to reduce subjectivities. An example of this is the application in the organization of the case backlog, which happens through the identification of whether a given case falls within the jurisdiction of a specific precedent. The tool can automatically forward the case to the staff member who has already worked on an analogous case, thus preventing the issuing of contradictory decisions, simply because similar cases were allocated to different staff members. This problem, which has always existed due to the extensive number of cases, directly impacts legal certainty, as the lack of consistency in judicial decisions represents one of the system's greatest challenges.

Furthermore, artificial intelligence may be successfully used in this manner as well as assisting the Judiciary Branch in achieving faster judgments. Thus, AI faces an old legal dilemma: the judge, in view of an excessive number of cases, must decide between judging with attention and detail, which will accumulate the backlog of cases, or judging quickly, which may increase the number of errors.

The structure and architecture of these tools allows for tackling both fronts simultaneously. As a result, it's possible to obtain a significant gain in effectiveness, with more sound, secure and reliable decisions, in addition to the contribution to legal certainty, while speeding up the judgment process.

[Interviewers] — What are the main risks in the use of artificial intelligence by the courts?

Rodrigo Canalli — The first, and most studied, is the matter of the so-called algorithmic biases. Machine learning models, which learn to recognize patterns based on data — vast amounts of data — often end up reproducing discriminatory biases that were already present. That data influences the continuity of prejudice and injustices. This is a great problem in the governance of data and the governance of artificial intelligence. Within legal matters we have many cases that include examples of compared law. Thus, when a machine is trained based on data that contains biases, it will learn exactly that, reproducing it and, eventually, even amplify these prejudices.

Another problem is the lack of transparency, or opacity, from artificial intelligence, also known as the “black boxes” phenomenon. It refers to the lack of knowledge concerning the parameters or criteria that these tools use to make decisions or reach a specific result. This problem may have two main sources. Primarily, the black boxes may result from questions related to intellectual property such as software protected by trade secrets, in which one cannot have access to the inner workings of that specific system. The second source is in the sheer complexity of these models, which is massive. Often, not even the engineers who develop the code that generated the algorithmic model can understand exactly how the model reaches specific results.

In Brazil, the issue of intellectual property in the Judiciary Branch has been less relevant because, until now, the Brazilian Judiciary has developed internal solutions. Most courts have their own IT teams. The Brazilian Judiciary has, historically, been a pioneer in the adoption of

digital solutions, and with artificial intelligence it was not different. Furthermore, the National Council of Justice (CNJ) has centralized some initiatives to establish a pattern but, in general, courts develop their own tools, share them, and make them available. In Brazil, we don't face many problems related to trade secrets or intellectual property, as they are usually public, open-source software available to all.

However, the problem of lack of transparency, due to the very complexity of the tools, persists. Due to their nature, the architecture of these systems, contributes to the difficulty in fully comprehending their inner workings, mainly when it comes to the software used to produce decisions or to assist in decision making process.

Through this, one of the major difficulties is the training of AI models, particularly the establishment of what's called the "ground truth" — it is the truth parameter against which the system will be evaluated. In the context of a court, defining this parameter might be extremely complicated, considering the conflicting precedents and the ambiguities inherent to human and legal language.

Herein we also have concerns related to the due legal process, since the Brazilian Constitution demands that all legal decisions should be public and reasoned. It's fundamental that the motives that led to a legal decision are transparent. If the result was produced by a software that, even though through sound reasoning, which may seem consistent, does not explain the argumentation on how that decision was achieved or which were the used criteria, we may find ourselves before a conflict between the adoption of these tools and the due legal process, which demands, amongst others, transparency.

Moreover, the legitimacy of the judicial decision itself may be compromised, as it possesses a discursive legitimacy, grounded in a rational activity conveyed through the reasoning provided. There are clear examples of this. Our legislation, such as the General Law of Data Protection (LGPD), does not mention artificial intelligence specifically, although it does encompass it. LGPD establishes legal parameters that demand the individual's access to the criteria used by tools in order to make decisions capable of impacting their rights or interests, including decisions made by automated systems.

However, significant questions remain. What exactly characterizes a fully automated decision? In which cases does an individual have the right to a review made by a human being? And, furthermore, what is considered a valid reasoning for these decisions? These

questions are still unanswered and are widely debated within the interactions between law and technology, especially in artificial intelligence.

[Interviewers] — What are the main ways of development of artificial intelligence tools in the Brazilian Judiciary?

Rodrigo Canalli — There are three main forms for the development of artificial intelligence tools in the Judiciary. The first one is through collaboration with universities, such as in the partnership between the Supreme Federal Court (STF) and the University of Brasilia for the development of Victor. The second way is the internal development by the court's own IT teams. As previously mentioned, both tools, RAFA 2030 and VitorIA, belong to the Supreme Court, and were entirely developed by the court's own teams. The third way is through a partnership with the private sector. The STF's newest AI, MARIA, was developed in partnership with EloGroup and Microsoft, using Meta's language model (LLM).

[Interviewers] — How does the National Justice Council currently regulate the use of artificial intelligence in the Judiciary Branch?

Rodrigo Canalli — Currently, the use of artificial intelligence in the Brazilian Judiciary Branch is mainly regulated by the Resolution n. 615, edited by the National Justice Council in March 2025. This normative updated and replaced the old Resolution n. 332/2020, including the most recent advances in the AI field, even the developments related to extensive language models and the generative artificial intelligence technologies.

The Resolution establishes fairly rigorous parameters for the development, governance and use of these tools in the Judiciary. The resolution adopts the premise that any AI solution must be aligned with principles such as ethics, transparency, safety and, above all, the protection of fundamental rights. This means that all systems must have clear mechanisms of human supervision, robust documentation, measures for mitigation of discriminatory biases and guarantees of transparency and auditability.

An important topic is that, even though this normative does not have binding effects for the Supreme Federal Court, it has strong persuasive effect, serving as a reference for the adoption of good practices in technological governance throughout the entire justice system.

Amidst the guiding principles of such regulation, I would highlight three pillars. The first one refers to non-discrimination, which demands previous homologation of models in order to

identify and eliminate possible biases. If these biases are irremovable, the use of the system is discontinued. The second pillar is transparency, which demands detailed documentation of the goals, risks, mitigation strategies and, mainly, the capacity of the system to offer auditable explanations of its results. And, finally, the third pillar is institutional control and auditability, which guarantee that the systems may be permanently revised and submitted to public scrutiny.

In addition to these principles, the Resolution imposes a series of specific prohibitions. As an example, it's absolutely prohibited that AI tools substitute legal reasoning or that they generate binding decisions automatically. There are also restrictions to the use of more sensitive technologies, such as facial recognition, which can only be adopted with direct authorization from the National Justice Council (CNJ). And there's a very rigorous requirement related to technical documentation and communication of new projects to the National Justice Council (CNJ) itself, including the obligation of depositing the models in the Sinapses repository.

Anyway, it's important to note that this is still a work in progress. The current regulation offers initial parameters, but the speed of technological evolution, especially in the field of AI, demands constant reflection. The challenge, in the end, is guaranteeing that these innovations will always be subordinate to the constitutional principles and the values that structure the Rule of Law.

[Interviewers] — What are the main artificial intelligence tools developed by the Supreme Federal Court, how do they function and how do they contribute to the jurisdictional activity?

Rodrigo Canalli — The first tool developed was Victor, launched in 2018, during the presidency of Justice Cármen Lúcia. It was made through a partnership with the University of Brasília (UnB). Today, Victor is kept, updated and enhanced fully by the court's own IT teams. Its' main function is classifying cases according to general repercussion themes (STF's system of classification of binding precedents), with an accuracy average **above** 80%. It does not have the final say in any decision, and all suggestions made by Victor are reviewed by two of the court's teams.

The second tool developed by STF was RAFA 2030, the name is an acronym for Artificial Networks Focused on UN's 2030 Agenda. RAFA classifies the court's cases according to the

UN's Sustainable Development Goals (SDG), with a precision superior to 90%. Just like Victor, RAFA does not have autonomy to decide, being limited to suggesting classifications that are always confirmed by civil servants.

Later, during the presidency of Justice Rosa Weber, the tool VitorIA was developed, in which I had the honor of being part of the team responsible for the tool's idealization, development and launching. VitorIA is a tool meant to group cases through textual similarity. It analyses the court's cases backlog by class, arrival in the Court or jurisdictional unity, then it reads some selected procedural documents and groups the cases that discuss similar themes. This functionality significantly speeds up the court's jurisdictional activity.

VitorIA also allows the court to identify litigation waves and new general repercussion themes in advance, something that used to be diluted between cabinets and civil servants. Since its implementation, some themes submitted to virtual judgment already had the general repercussion identified through VitorIA's analyses.

Additionally, the tool offers other functionalities. For example, when it identifies and creates a group of cases on a specific theme, all new cases related to that topic are automatically allocated. Similarly, VitorIA assists civil servants and judges in the identification of cases similar to the ones they are currently analyzing, allowing the use of previously made analyses, which contributes for a greater efficiency in handling the caseload. Currently, VitorIA works with extraordinary appeals, appeals against decisions on extraordinary appeals, and complaints, which are the highest volume case types in the court. The expansion of its scope to include *habeas corpus* is in the developmental stage.

It is important to note that, like Victor and RAFA 2030, VitorIA does not make any decisions independently, limiting itself to offering groupings and support for civil servants and judges.

In December 2024, MARIA (Module for AI-Assisted Drafting) was launched. It is the Supreme Federal Court's first tool which incorporates generative artificial intelligence models into its legal and administrative support routines, expanding the possibilities for intelligent automation of written tasks within the Court. The tool assists in drafting summaries of judicial opinions, reports for appeal cases, and the preliminary analysis of cases in the Complaint class.

[Interviewers] — What are the impacts of biases in artificial intelligence models?

Rodrigo Canalli — The matter of biases is present in the discussions about artificial intelligence models, especially because these models learn from our language, as they are, mainly, language models. They process, understand, and produce human language, reflecting the vices that we, as a society, insert in them. Language is, after all, what we use to communicate.

To deal with those biases, it's necessary to adopt different approaches. One of them is the composition of diverse teams. Teams that work with artificial intelligence must be diverse, in such a way that they can't be composed only of "traditional white males." Even though this is a valid point of view, it is also limited and frequently ignores other groups' perceptions that are essential for the construction of tools of this nature.

A clear example of this happened at the beginning of the development of digital photographic cameras, which used face identification. These cameras identified the faces of white men very well, presenting a slightly inferior performance when it came to white women's faces, and identified the faces of black people poorly — the worst results were with black women. That happened because, in the teams that developed these tools, there wasn't a concern with diversity.

Furthermore, a truthful saying is that the most important decisions on technology are not only technical, but also political. Technology is never neutral; it has a role that reflects the way it is actively inserted in our society, influencing economic and social models and every type of connection.

That's why a whole discussion on the regulation of technologies emerges. There's no neutrality in this context; all choices are political. When we choose to not regulate something, we are, essentially, taking a political decision, not just a technical one.

Another example that illustrates this matter is feminism. I remember having asked ChatGPT to create an image of a black judge. All images generated presented a woman in sensual positions, with accentuated necklines. That happened because the tool learned, based on the language and data it was trained with, to produce these types of results. These prejudices are rooted in our language and end up being transferred to these technologies. The artificial intelligence tools must incorporate perspectives such as feminist constitutionalism, black constitutionalism, and anticolonial constitutionalism.

From the geopolitical perspective, the manners these tools are structured and economically distributed tend to reproduce a new type of digital colonialism. Even though these artificial intelligence tools are developed by big companies from the northern hemisphere, its functioning depends on the training made by many poorly paid people in the southern hemisphere, mainly in Africa and Latin America. This scenario reflects a financial disparity, which is part of the current technological context.

[Interviewers] — What are the perspectives for the use of high-risk artificial intelligence tools in the Judiciary in the next few decades?

Rodrigo Canalli — I believe that we are heading towards an increasingly higher usage of tools that act directly in decision making activities, which would be classified as high-risk. The tools will need to be adapted to fit the National Justice Council's normative in order to reduce risks. However, I agree with the global tendency of the expansion of artificial intelligence tools directed towards decision making.

I don't see this as something exclusively risky. Judgments made by people also present risks. Thus, even though I observe the growing influence of AI in decision making, I believe that they should not have the final word, at least not in the near future. Even so, AI's strong influence in decision making processes' will be unavoidable.

ACKNOWLEDGMENT

[Interviewers] — Professor, in name of the Revista de Egressos e Acadêmicos de Direito and the Centro Brasileiro de Estudos Constitucionais, both from CEUB, we thank you immensely for your contribution and hope to see you soon.

Rodrigo Canalli — Thank you very much for the opportunity. I'm at your disposal.

Technical Details

Journal name: Revista de Egressos e Acadêmicos de Direito do Centro Universitário de Brasília (READ CEUB)

Volume and year: V. 1, N. 1 – 2026

Interview title: Interview with Professor Rodrigo Canalli on Artificial Intelligence in the Brazilian Judicial Branch

Interviewee name: Rodrigo Lobo Canalli

Interviewers' names: Christine Oliveira Peter da Silva, Ian Ferrare Meier, Luísa Cristina Vasconcelos Marimon Álvares e Lucas de Pádua Carvalho Mendes Lima

Translators and translation reviewers' names: Ester Azevedo Lira e Pedro Galuppo Inácio

E-ISSN: xxxxx