

Judge Hercules And Laplace's Demon: Philosophical and Practical Limits of Predictive Artificial Intelligence in Judicial Decision-Making

Diógenes Vicente Hassan Ribeiro¹, Pedro Ramos Lima²

¹Undergraduate Law Program at Unilasalle

²La Salle University

ABSTRACT: Recent advances in artificial intelligence have renewed interest in predictive models of judicial decision-making, raising questions about determinism, legal rationality, and the possibility of overcoming uncertainty in law. This article examines whether contemporary AI-based predictive systems revive idealized models of legal determinacy, drawing on Dworkin's Judge Hercules and Laplace's Demon as theoretical reference points. Through a conceptual and comparative analysis grounded in legal theory, philosophy of science, and socio-legal research, the article traces the historical development of predictive models and evaluates their cognitive, epistemological, and systemic limits. It argues that predictive accuracy, even when surpassing human expertise in specific tasks, does not transform law into a deterministic system. Instead, predictive models generate probabilistic expectations rooted in past decisions and conditioned by temporal, institutional, and social dynamics. By situating artificial intelligence within the structural uncertainty of law, the article shows that AI may enhance consistency while leaving indeterminacy intact, unaltered overall.

KEYWORDS: Artificial intelligence; judicial decision-making; legal uncertainty; predictive models

1. INTRODUCTION

Contemporary legal research employing artificial intelligence (AI), particularly in the development of predictive models of judicial decision-making, has produced results that are both technically impressive and theoretically provocative. Recent studies demonstrate not only the capacity of such models to process and extract patterns from millions of judicial decisions, but also their ability to predict case outcomes with levels of accuracy that, in specific and well-defined tasks, may exceed those of human legal experts. A paradigmatic example is provided by Menezes-Neto and Clementino (2022, p. 17), whose models analyzed more than two million decisions from the Brazilian Federal Small Claims Courts and achieved higher predictive accuracy in appellate outcomes than members of the federal judiciary. These findings have reinforced the perception that AI systems, at least in terms of information processing and pattern recognition, operate at a scale that appears "superhuman" when compared to human cognitive capacities.

This perception, however, immediately raises a classical theoretical problem: whether advances in prediction signal a renewed possibility of determinism. In the philosophy of science, this problem is traditionally illustrated by Laplace's Demon, a hypothetical intellect capable of knowing all forces and positions in nature and, on that basis, determining the future and the past with complete certainty. Transposed to the legal domain, the question becomes whether sufficiently powerful predictive models could dissolve legal uncertainty by rendering judicial outcomes fully foreseeable. The contemporary enthusiasm surrounding AI-based legal prediction thus revives, in a new technical form, an old epistemological tension between determinacy and uncertainty.

Legal theory has long grappled with a parallel aspiration. Ronald Dworkin's figure of Judge Hercules famously represents an idealized response to legal indeterminacy: a jurist endowed with unlimited time, patience, and cognitive capacity, capable of integrating all relevant legal materials into a coherent political theory and producing a single right answer for every hard case (Dworkin, 2010, p. 165). Although explicitly conceived as a theoretical construct rather than a practical model, Judge Hercules embodies a rational ideal structurally analogous to Laplace's Demon. The recent success of predictive technologies raises the question of whether artificial intelligence might approximate, in functional terms, the role attributed to these idealized figures, thereby transforming the relationship between law, prediction, and uncertainty.

This article addresses that question through a theoretical and conceptual analysis grounded in a selective literature review and comparative examination of key contributions in legal theory, philosophy of science, and socio-legal studies. The analysis mobilizes representative works from four main lines of research: early quantitative and statistical models of judicial behavior; machine learning and natural language processing approaches to legal texts; neural network and transformer-based architectures applied to legal prediction; and empirical socio-legal studies focused on concrete judicial systems, with particular attention to recent

Judge Hercules And Laplace's Demon: Philosophical and Practical Limits of Predictive Artificial Intelligence in Judicial Decision-Making

research conducted in Brazil. These works are selected not for exhaustiveness, but for their capacity to illuminate decisive moments in the historical and conceptual evolution of predictive modeling in law.

Rather than evaluating predictive accuracy as an end in itself, the article focuses on the conceptual assumptions underlying claims of determinacy, objectivity, and completeness in AI-assisted adjudication. It examines how predictive models implicitly reproduce idealized notions of legal rationality and explores the extent to which such notions remain compatible with the way law actually operates as a social practice. The discussion therefore shifts attention from technical performance to the structural conditions under which legal decisions are produced.

Finally, the article incorporates insights from behavioral psychology and theories of bounded rationality to address a dimension often overlooked in techno-legal debates: the cognitive limits of decision-makers. Contrary to the premise of unlimited information accumulation that underpins idealized models of adjudication, human decision-making depends on selective attention and heuristic processing to avoid cognitive overload—a constraint that, in different terms, also affects computational systems. By adopting a systemic perspective, the article argues that artificial intelligence should be understood not as a mechanism for eliminating legal uncertainty, but as a tool that operates within an inherently uncertain and temporally structured legal system, in which consistency and predictability remain necessarily provisional.

2. JUDGE HERCULES, LAPLACE'S DEMON, AND LEGAL DETERMINISM

The science of predictive modeling and Ronald Dworkin's theoretical construct of Judge Hercules share a significant structural affinity: both propose models of law aimed at overcoming uncertainty by rendering legal outcomes intelligible and, to some extent, predictable. On the one hand, predictive models grounded in mathematics, algorithms, and computational technologies seek to transform past judicial decisions into datasets from which patterns can be extracted and future outcomes estimated. On the other hand, Dworkin's Judge Hercules embodies a normative and theoretical model centered on a single, comprehensive legal theory capable of yielding one correct answer to every legal dispute. Despite their different foundations, both approaches implicitly engage with the aspiration to reduce legal indeterminacy by positing a form of rational completeness.

The divergence between these models lies primarily in their points of departure and methodological orientation. Predictive models are empirically driven: they investigate past judicial decisions in order to identify statistically recurrent patterns that may inform expectations about future rulings. Judge Hercules, by contrast, is an abstract and counterfactual figure designed to articulate a theory of law that is, in principle, independent of the contingent practices of courts and tribunals. Indeed, Dworkin's construct often stands in critical tension with judicial practice, insofar as it challenges positivist accounts that privilege institutional behavior over normative coherence. The similarity between the two models, therefore, should not obscure their distinct ambitions: one seeks empirical predictability, while the other seeks normative determinacy.

Judge Hercules is introduced by Dworkin as a theoretical exercise intended to clarify the meaning of the "right answer" thesis. Hercules is described as a jurist endowed with "superhuman skill, learning, patience, and acumen" (Dworkin, 2010, p. 165), capable of knowing and understanding everything that can be known about a given legal system. On the basis of this exhaustive knowledge, Hercules would be able to construct a comprehensive political theory capable of justifying the constitution as a whole (Dworkin, 2010, p. 166) and, from that standpoint, to deliver a correct and conclusive answer to any legal problem, provided that such an answer exists within the system. Importantly, this model does not rely on mechanical rule application, but on an integrated interpretation of legal principles, institutional history, and political morality.

Dworkin's philosophical model thus operates as a reaction to legal uncertainty. It proposes a morally grounded legal standard that, in the absence of an explicit rule, supplies a binding criterion for judicial decision-making (Neves, 2016, p. 47). To sustain this standard, Dworkin invokes a superhuman judge, an ideal or transcendental subject, who possesses complete knowledge of a nation's law and is capable of articulating a legal theory sufficiently coherent and informed to identify the relevant principles in any legal controversy. Through this theory, Hercules would make it possible to arrive at a single correct answer for every case submitted to adjudication (Neves, 2016, p. 47). The figure of Hercules, therefore, should not be understood as a realistic model of judging, but as a heuristic device designed to illuminate the normative aspirations of a constructivist theory of law.

Within Dworkin's constructivist framework, the identification of the right answer in adjudication requires judges to orient themselves toward the political ideal of the community. This ideal is not derived from the personal moral convictions of individual judges, but from a coherent and structured whole grounded in the community's political morality, the tradition of institutional decisions, and the jurisprudence of courts and tribunals (Silva and Tizzo, 2023, p. 270). The complexity of this interpretive task is such that it presupposes an agent capable of integrating vast and heterogeneous normative materials into a single justificatory framework. It is precisely this extreme cognitive and interpretive demand that explains the necessity, within Dworkin's theory, of an idealized figure such as Judge Hercules, one whose superhuman capacities render feasible a task that would otherwise exceed the limits of ordinary judicial reasoning.

Dworkin's theoretical construct of Judge Hercules bears a significant structural resemblance to the classical thought experiment proposed by Pierre-Simon Laplace to explain what later became known as "Laplace's Demon." In the context of classical

Judge Hercules And Laplace's Demon: Philosophical and Practical Limits of Predictive Artificial Intelligence in Judicial Decision-Making

physics, Laplace imagined an intellect vast enough to know all the forces governing nature and the precise position of every object composing it, down to its smallest particles. If such an intellect existed, it would be capable of grasping, within a single formula, the motion of all bodies in the universe; by solving that formula, nothing would remain uncertain, and the future, as well as the past, would be fully present to its understanding (Laplace, 1902, p. 4). The demon thus represents the strongest expression of determinism: complete knowledge of initial conditions, combined with universal laws, would render the evolution of the world entirely calculable.

Within physics, Laplace's demon presupposes an absolute intelligence capable of identifying all the fundamental components of nature and the forces acting between them. By knowing the totality of initial conditions, such an intellect could formulate the differential equations that describe the laws of nature and govern physical dynamics. Equipped with unlimited computational capacity, it could then solve those equations and thereby eliminate all uncertainty regarding nature, whether past, present, or future (Frigg *et al.*, 2014, pp. 32–33). Determinism, in this sense, refers to the idea that the state of a system at any given moment is fully determined by its prior state in accordance with physical laws, such that, in principle, all future states are predictable once the initial conditions are known (Cattani *et al.*, 2016, p. 2).

A similar determinist aspiration can be identified in both Laplace's exercise in physics and Dworkin's exercise in legal theory, despite their distinct domains and methodologies. In each case, the system, whether physical or legal, is assumed to be intelligible through a comprehensive rational framework capable of accounting for all relevant elements. For Laplace, this framework takes the form of differential equations governing nature; for Dworkin, it appears as a general political theory capable of justifying the law as a coherent whole. Judge Hercules, by virtue of his exhaustive knowledge of legal materials and principles, would be able to integrate statutes, precedents, and constitutional norms into a single interpretive scheme, producing determinate answers to legal disputes without remainder.

It is crucial, however, to recognize that the analogy between Judge Hercules and Laplace's demon is structural rather than ontological. While Laplace's demon operates within a causal model of the universe, Dworkin's Hercules functions within a normative and interpretive framework. The determinism at stake in Dworkin's theory does not concern causal necessity, but the possibility of normative completeness: the idea that, for every legal controversy, there exists a uniquely correct answer that can be identified through sufficiently comprehensive interpretation. Hercules does not mechanically derive outcomes from rules; instead, he constructs the best possible interpretation of the legal system by reconciling its institutional history with the moral principles that justify it.

Nevertheless, both figures share a common ambition: the elimination of uncertainty through rational integration. In the case of Laplace, uncertainty disappears once the laws of nature and initial conditions are fully known. In the case of Judge Hercules, uncertainty vanishes once the legal system is interpreted as a coherent whole under a single justificatory theory. This parallel illuminates why contemporary enthusiasm for predictive models of judicial decision-making so readily evokes comparisons with idealized forms of legal rationality. By promising increasingly accurate forecasts of judicial outcomes based on the systematic analysis of past decisions, such models appear to revive, in technical form, the aspiration to overcome indeterminacy through comprehensive knowledge, an aspiration that encounters decisive limits when confronted with the temporal, contingent, and socially embedded nature of law.

A central question that challenges deterministic models of explanation is whether it is in fact necessary to derive and solve a system's fundamental equations in order to understand its behavior. In many domains of knowledge, meaningful and robust explanations have been achieved without access to ultimate foundations, through the use of partial, approximate, and historically situated models. Biology offers a paradigmatic example. Darwin's theory of evolution by natural selection was formulated without knowledge of genetics or of the molecular mechanisms underlying inheritance. Yet it provided a powerful explanatory framework by identifying a mechanism of variation and selection capable of accounting for biological change over time. Rather than eliminating uncertainty, Darwin's theory institutionalized it as a condition of explanation, treating evolution as an open-ended and contingent process rather than as the realization of a closed, fully determinate model.

Even Laplace himself, in formulating the image of an all-knowing intelligence, did not claim that such an intellect was attainable. On the contrary, the figure of Laplace's demon served primarily to emphasize the remarkable achievements of human astronomy while acknowledging that humanity remains "infinitely removed" (Laplace, 1902, p. 4) from any intelligence capable of revealing the complete truth of all things. This recognition of epistemic distance becomes even more pronounced in modern physics, which began to question the core assumptions of classical determinism. With the advent of relativity theory, the inductive method itself was problematized, and the idea that empirical observation could mechanically lead to fundamental laws was explicitly rejected (Einstein, 1936, p. 365). Importantly, this shift does not imply a denial of regularity in nature, but rather a rejection of the belief that such regularity can be captured by a final and complete theoretical framework.

Einstein's reflections on theoretical physics underscore this point. For him, the fundamental concepts of physics are not passively derived from experience, but actively constructed through what he described as "free inventions of the human mind" (Einstein, 1936, p. 381), whose justification lies not in their origin but in their explanatory and predictive utility. Modern physics thus operates in a state of permanent conceptual revision, in which theories are continuously tested, reformulated, and sometimes

Judge Hercules And Laplace's Demon: Philosophical and Practical Limits of Predictive Artificial Intelligence in Judicial Decision-Making

abandoned in light of new problems. This indeterminism is methodological and epistemological rather than ontological: it affirms the provisional character of scientific explanations without abandoning the search for coherence or consistency.

A similar shift can be observed in legal theory when law is approached as a complex social phenomenon rather than as a closed normative system. Attempts to model law as a fully determinate order governed by comprehensive theoretical principles face difficulties analogous to those encountered by classical physics when confronted with increasingly complex phenomena. As legal systems evolve through legislation, adjudication, and social change, their normative structures are constantly reconfigured, rendering implausible the expectation of a single theory capable of yielding definitive answers to all legal questions. In this sense, the challenge to determinism in law does not entail relativism, but rather a recognition of contingency and revision as structural features of legal reasoning.

This perspective casts doubt on the explanatory sufficiency of deterministic theoretical models such as Judge Hercules or Laplace's demon, even when understood strictly as heuristic devices. Both figures presuppose systems that can be rendered fully intelligible through comprehensive rational integration. Yet the systems they seek to illuminate, physical reality and law, are characterized by temporalization, adaptation, and continuous transformation. As Einstein observed in response to the breakdown of classical theories of light and electricity, adherence to a single, all-encompassing framework may obstruct critical inquiry rather than advance it (Einstein, 1936, p. 381). Applied to legal reasoning, this insight suggests that the pursuit of a unique correct answer grounded in a complete theory risks obscuring the dynamic and socially embedded character of law. Rather than deriving legal outcomes from foundational equations, understanding law requires engaging with its evolving structures, its contingent operations, and the limits inherent in any attempt at theoretical closure.

3. JUDGE HERCULES, INFORMATION ACCUMULATION, AND COGNITIVE OVERLOAD

Beyond the theoretical plausibility of Judge Hercules as a model for legal theory, its practical impossibility becomes evident when the construct is contrasted with how human cognition actually operates. Dworkin's ideal presupposes a mode of reasoning based on the progressive accumulation of information, culminating in a comprehensive political theory capable of justifying the legal system as a whole. Human decision-making, however, does not function through unlimited aggregation of data. On the contrary, it operates through mechanisms of selection, simplification, and prioritization that bear greater resemblance to contemporary scientific practices centered on testing the internal consistency and practical adequacy of theoretical formulations than to the classical aspiration of deriving a single, definitive answer from complete knowledge.

The figure of Judge Hercules embodies the assumption, widely shared in common-sense reasoning, that better decisions necessarily follow from having access to more information. According to this view, decision-making errors are typically attributed to informational deficits: the decision-maker did not know enough, did not have access to all relevant facts, or failed to consider certain variables. Much less attention is paid to the opposite possibility, namely that error may result from an excess of information. Human cognitive capacities are limited, and when the volume of available data exceeds the ability to process it meaningfully, decision-makers become vulnerable to what has been described as cognitive overload. Under such conditions, the presence of additional information does not improve judgment; it impairs it (Wilkins, 1973, p. 190).

Dworkin's model relies precisely on the negation of this limitation. Judge Hercules is conceived as a jurist who has at his disposal all the information necessary to construct a theoretical model of law capable of encompassing every possible circumstance and yielding a single correct answer to each legal problem (Dworkin, 2010, p. 166). At an abstract level, this model appears compelling: if one could process all relevant information without constraint, the resulting decision would presumably satisfy the internal requirements of coherence and justification demanded by the theory. The difficulty arises when this ideal is transposed to the level of actual adjudication. No matter how skilled or experienced, human judges are capable of processing only limited quantities of information at any given time. Their reasoning depends not on exhaustive accumulation, but on selective attention and heuristic strategies designed to manage complexity (Wilkins, 1973, p. 190).

This practical limitation has important implications for legal theory. It suggests that the aspiration to eliminate uncertainty through comprehensive informational control is not merely unattainable, but misguided. Legal decision-making does not improve indefinitely as information increases; rather, it requires institutional and cognitive mechanisms that reduce complexity to manageable levels. In this sense, the practical critique of Judge Hercules converges with the broader epistemological critique developed earlier: just as modern science abandoned the expectation of complete theoretical closure, legal reasoning must be understood as operating under conditions of bounded rationality, structural uncertainty, and ongoing revision. The challenge for legal theory, therefore, is not to imagine an idealized decision-maker immune to cognitive limits, but to explain how law functions and stabilizes expectations precisely because such limits cannot be overcome.

Rutkowski and Saunders (2019, p. 5) employ a useful metaphor to explain how the human brain processes information, comparing it to a blender. In a blender, solid pieces are placed inside and processed by activating different power levels. If the pieces are too large or if a finer consistency is desired, greater power is required. However, if too many pieces are added, the blender fails to process them adequately, resulting either in an unsatisfactory mixture or, in extreme cases, in overflow and loss of the material

Judge Hercules And Laplace's Demon: Philosophical and Practical Limits of Predictive Artificial Intelligence in Judicial Decision-Making

intended to be preserved. This metaphor captures an essential feature of human cognition: information must be processed in limited and manageable units.

Similarly, the human brain processes information in fragments rather than as an undifferentiated whole. Under ordinary conditions, when information is relatively simple or limited in quantity, it can be processed without difficulty. Yet there are situations in which the available cognitive resources are insufficient to handle the volume or complexity of incoming information. In such cases, individuals must either mobilize higher levels of cognitive effort or adapt by processing information less thoroughly. Both strategies increase the likelihood of error, reduce informational integrity, and impair decision-making, particularly because sustained access to higher levels of cognitive processing leads to psychological exhaustion and depletion of mental resources (Rutkowski and Saunders, 2019, p. 6).

Importantly, cognitive overload is not restricted to non-experts. Specialists are generally able to process information more efficiently and, when confronted with the same informational load as non-specialists, tend to expend less cognitive effort to perform a given task (Rutkowski and Saunders, 2019, p. 10). Nonetheless, experts remain vulnerable to overload, especially in contexts of technological change that require them to unlearn previously internalized procedures in order to adopt more efficient ones (Rutkowski and Saunders, 2019, p. 11). Learning new procedures places experts in a position similar to that of novices, requiring access to higher cognitive resources. This condition is particularly characteristic of contemporary societies, in which rapid technological innovation continuously increases both the quantity of information available and the demands placed on cognitive processing.

The problem of cognitive overload, however, is not exclusive to the so-called information age. It reflects a more fundamental aspect of the human relationship with information across history. Humans have never been capable of cataloging, analyzing, or storing all the information that passes through their senses. The world beyond individual cognition—extending beyond what can be directly processed by the mind—has always been infinitely complex, long before the discovery of subatomic particles and their associated indeterminacies (Watson, 2019, p. 6). Cognitive limitation, therefore, is not a technological anomaly, but a structural condition of human existence.

According to Watson (2019, p. 7), culture functions as a mediating thread that helps manage the tension between the cacophony of the world and the limited capacity to perceive and interpret it. Cultural frameworks provide filters and pathways that guide attention toward information deemed relevant in a given context. This selective process is inherently contingent and closely resembles what Luhmann describes as “order from noise,” a concept used to explain how communication emerges through the selective attribution of meaning to an excess of information (Luhmann, 2013, p. 84). In this sense, the challenge of information processing extends beyond individual cognition and becomes a social problem.

From a systemic perspective, when the organization of human cognition is extrapolated to social organizations, decision-making processes likewise operate through mechanisms of selection designed to prevent informational overload. This stands in contrast to the widespread assumption that the best decision is necessarily the most informed one. Decision-making within organizations is not a purely rational process in which all available information is processed exhaustively. Rather, it depends on the exclusion, simplification, and prioritization of information. To understand how decisions are produced by social organizations, it is therefore insufficient to focus solely on rational choice; it is necessary to examine which information is ignored, which is selected, and how such selections are structured. As Luhmann (2018, p. 371) emphasizes, this is a central problem of information dependency in decision-making processes and highlights the limits inherent in any attempt to equate better decisions with greater quantities of information.

Judge Hercules is a theoretical model that systematically disregards the phenomenon of cognitive overload and its impact on decision-making processes (Levy, 2022, p. 160). The idealization of a judge capable of considering all factors relevant to legal rationality, comprehending every applicable rule and principle involved in a case, as if assembling an exhaustive “shopping list” of legal reasons, fails to account for the limits of human memory and attention. Cognitive psychology has long demonstrated that human agents are constrained in their ability to store, retrieve, and process information. When confronted with an excessive volume of information, decision-makers tend either to ignore relevant factors or, in extreme cases, to suspend decision-making altogether, overwhelmed by the complexity of the task.

The determinism implicit in the figure of Judge Hercules further compounds this difficulty. By presupposing the existence of a single correct answer or a perfect political-legal theory capable of justifying the constitution as a coherent whole, Dworkin's model attributes to law a level of internal consistency in which rights, rules, and principles are applied uniformly to ensure that all cases are decided in a fair and equal manner. Levy (2022, p. 161) characterizes this assumption as a form of “metaphysical independence of value,” according to which legal outcomes are determined independently of the contingent conditions under which decisions are actually made. While this premise may serve a valuable heuristic function in theoretical debates, it finds no practical equivalent in the everyday operation of courts and tribunals. In practice, judicial decision-making is less concerned with uncovering a pre-existing correct answer than with determining what constitutes the best or most appropriate outcome in a given case, a task that inevitably shifts attention toward procedures capable of managing uncertainty and optimizing results.

Judge Hercules And Laplace's Demon: Philosophical and Practical Limits of Predictive Artificial Intelligence in Judicial Decision-Making

Even if one were to idealize a model of judicial decision-making inspired by Judge Hercules, the Dworkinian framework would offer little guidance for human judges operating under cognitive constraints. The assumption of superhuman knowledge regarding the multiple elements that compose legal rationality is incompatible with the reality of adjudication, where judges must focus on a limited set of factors and actively prioritize which pieces of information are most relevant to resolving a dispute. As Levy (2022, p. 163) argues, a judge who attempted to follow Hercules' recommendations by incorporating ever more information into the decision-making process would likely produce less satisfactory outcomes, precisely because excessive information undermines, rather than enhances, effective judgment.

From this perspective, the ideal model for human decision-making is not one that maximizes the amount of information available to the decision-maker, but one that explicitly accounts for cognitive overload by structuring and constraining the decision environment. Such a model seeks to identify which factors deserve priority and which information can be safely excluded, functioning as a filter rather than as an accumulator of data. Although Levy's analysis is grounded in psychological research on decision-making, it resonates with broader methodological shifts observed in modern science, where emphasis has moved away from the pursuit of all-encompassing foundational theories toward the development of coherent, context-sensitive models valued for their internal consistency and practical utility. Applied to law, this insight underscores the limits of idealized determinism and reinforces the need to understand adjudication as a selective, bounded, and institutionally mediated process.

4. ARTIFICIAL INTELLIGENCE, CONSISTENCY, AND LEGAL UNCERTAINTY

If the ideal model of decision-making is one that explicitly takes cognitive overload into account and therefore requires mechanisms to determine which factors and pieces of information should be prioritized (Levy, 2022, p. 163), the incorporation of computers and artificial intelligence appears, at first glance, as a way to circumvent human cognitive limits. Unlike human decision-makers, machines are not constrained by attention, fatigue, or memory decay in the same manner. From the earliest studies comparing human cognition and machine processing, computers were already recognized as possessing vastly superior capacities for information processing, even if their memory was initially limited by hardware configurations (Loftus and Loftus, 1976, p. 128). In contemporary contexts, particularly with the development of machine learning and natural language processing techniques, such limitations have been substantially mitigated. Current systems are capable of processing massive textual corpora, including databases containing thousands of books (Radford *et al.*, 2018, p. 4) and millions of judicial decisions (Menezes-Neto; Clementino, 2022, p. 13).

In the legal domain, these technological advances have enabled artificial intelligence systems not only to analyze large volumes of judicial decisions across multiple jurisdictions, such as the European Court of Human Rights (Aletras *et al.*, 2016), courts in China (Li *et al.*, 2019), and Brazil (Menezes-Neto and Clementino, 2022), but also to predict judicial outcomes with levels of accuracy that, in specific contexts, exceed those of human legal experts. The objectives of these predictive models closely resemble those associated with the theoretical figure of Judge Hercules: the identification of a correct outcome for judicial decisions through the absorption of uncertainty into a framework that promises greater stability and consistency in law (Menezes-Neto and Clementino, 2022, p. 17). By transforming past decisions into structured datasets and extracting statistically significant patterns, such systems seek to reduce unpredictability in adjudication.

For Menezes-Neto and Clementino (2022, p. 2), the pursuit of consistency and stability in law is directly connected to the practical needs of lawyers and their clients, who seek to minimize surprises and manage risks when pursuing favorable judicial outcomes. From this perspective, computational tools function as instruments for expectation management, allowing legal actors to anticipate likely results and adjust strategies accordingly. It is important to note that this concern is not new to legal practice. Long before the advent of artificial intelligence, the legal field developed various techniques and informal predictive models, ranging from jurisprudential analysis to experiential heuristics, aimed at forecasting judicial behavior.

One of the earliest systematic attempts can be found in Kort's pioneering study of the United States Supreme Court. Treating judicial decisions as "highly uncertain" phenomena, Kort (1957, p. 1) employed qualitative mathematical methods to construct a predictive model based on numerical representations of factors recurrently present in decided cases. His approach combined variables such as the severity of the offense, defendants' psychological conditions, and procedural irregularities, most notably the absence of legal representation, in order to identify critical points at which judicial convictions were maintained or reversed. Applied retrospectively to resolved cases, this mathematical formulation was then projected prospectively as a model for anticipating future decisions (Kort, 1957, p. 12). Although rudimentary by contemporary standards, Kort's work established a foundational intuition: judicial behavior, despite its apparent unpredictability, might display identifiable regularities.

Following Kort's initial contribution, a wide range of approaches emerged, employing qualitative analysis, probability theory, and increasingly sophisticated mathematical techniques. Ulmer (1963) explored probabilistic models of judicial voting, while Keown (1980, p. 852) introduced catastrophe theory to capture discontinuities and abrupt shifts in judicial outcomes. Subsequent studies demonstrated that, under certain conditions, judicial decisions could be predicted with relatively high levels of accuracy (Segal, 1984, p. 896), and that judges' preferences exhibit systematic, though context-dependent, variation over time (Lauderdale

Judge Hercules And Laplace's Demon: Philosophical and Practical Limits of Predictive Artificial Intelligence in Judicial Decision-Making

and Clark, 2012, p. 863). Despite methodological differences, these models shared a common motivation: dissatisfaction with the perceived unpredictability of judicial decision-making and a desire to impose analytical structure on judicial behavior.

A significant methodological shift occurred with the incorporation of machine learning and automated textual analysis into predictive research. Rather than relying primarily on predefined variables selected by researchers, prediction increasingly centered on the classification and extraction of information directly from legal texts (Zhong *et al.*, 2018, p. 3541). In this context, machine learning can be understood as a statistical approach in which algorithms learn patterns from data that are iteratively fed back into the system as inputs (Sumathi and Janani, 2020, p. 3). Aletras *et al.* (2016) provided a landmark application of this approach by developing a predictive model for decisions of the European Court of Human Rights. Using binary classification, their model took textual features extracted from case documents as inputs and generated outputs indicating whether a violation of a specific article of the European Convention on Human Rights would be found.

Despite outperforming earlier non-machine learning approaches, models such as that of Aletras *et al.* (2016) faced important limitations. They relied on relatively superficial textual parameters and required substantial manual intervention by researchers to determine which factors would be included in the analysis. This dependence on human preselection not only demanded significant effort but also limited the models' capacity for generalization, making replication across different legal contexts difficult (Zhong *et al.*, 2018, p. 3541). In response, predictive research began to integrate neural networks capable of learning representations of legal knowledge directly from data (Zhong *et al.*, 2018, p. 3542). An influential example is provided by Luo *et al.* (2017), who proposed an attention-based neural network that jointly modeled outcome prediction and textual information extraction in a unified framework applied to judicial decisions in China.

This line of development continued rapidly. Although early neural network models achieved promising results, they were soon surpassed by architectures capable of delivering higher accuracy with reduced computational complexity. The introduction of the Transformer architecture (Vaswani *et al.*, 2017) marked a decisive turning point. Transformers enabled more efficient processing of long textual sequences and quickly became the foundation for widely adopted language models, including systems used in commercial applications and research contexts. Subsequent models, such as BERT (Devlin *et al.*, 2018), further refined these techniques and were incorporated into large-scale predictive systems. In the legal domain, such advances made it possible to process datasets comprising millions of judicial decisions and to generate predictions that, in narrowly defined tasks, surpassed the performance of human specialists (Menezes-Neto and Clementino, 2022, pp. 13, 17). What distinguishes contemporary artificial intelligence is not the emergence of prediction itself, but the scale, speed, and apparent objectivity with which such predictions can now be generated. This shift raises fundamental questions about the relationship between technological consistency and the enduring uncertainty that characterizes law.

5. THE THEORETICAL AND PRACTICAL LIMITS OF PREDICTIVE MODELS: UNCERTAINTY, INCONSISTENCY, AND THE TEMPORALIZATION OF LAW

The long-standing concern with predicting judicial decisions, explored across mathematics, statistics, and computer science, closely parallels a central aspiration of legal theory: the desire to describe law with precision and to reduce uncertainty in adjudication. Across different theoretical formulations, this aspiration has taken the form of a search for coherence, consistency, and determinacy. Dworkin's model of Judge Hercules exemplifies this orientation by proposing that even the most controversial cases admit a single correct answer, identifiable through a comprehensive political theory capable of justifying the legal order as a whole (Dworkin, 2010, pp. 164–165). Predictive models of judicial decisions emerge from a similar dissatisfaction with the apparent opacity and variability of judicial outcomes, already observed in early attempts to model Supreme Court decisions mathematically (Kort, 1957, p. 2).

This shared aspiration becomes particularly visible in so-called hard cases, in which legal texts and precedents fail to yield unequivocal solutions. In Dworkin's framework, such cases motivate the construction of an idealized rationality capable of absorbing uncertainty into coherence, rejecting the idea that indeterminacy necessarily leads to discretion or arbitrariness (Dworkin, 2010, p. 127). In predictive modeling, hard cases similarly motivate the search for regularities within large datasets of past decisions, with the aim of rendering future outcomes foreseeable. In both contexts, indeterminacy is approached as a problem to be solved rather than as a constitutive feature of law.

Contemporary predictive models, particularly those employing artificial intelligence, initially appear to vindicate this expectation. By processing vast corpora of judicial decisions and identifying statistically significant patterns, these systems can anticipate outcomes with high levels of accuracy, in some cases surpassing human legal experts in narrowly defined tasks (Menezes-Neto and Clementino, 2022, pp. 13, 16–17). Yet this apparent success simultaneously exposes their fundamental limitation. Because predictive models are trained on historical data, their capacity to foresee future decisions remains inseparable from the persistence of past decision-making patterns. As institutional practices, interpretive frameworks, and social expectations evolve, predictive accuracy inevitably degrades, requiring constant recalibration and revision.

Judge Hercules And Laplace's Demon: Philosophical and Practical Limits of Predictive Artificial Intelligence in Judicial Decision-Making

What emerges from both theoretical reflection and empirical research is that law has always contained an irreducible element of uncertainty. This uncertainty is neither accidental nor pathological; it is structurally linked to the way law operates within modern society. Legal decision-making unfolds within temporal dynamics that allow for variation, resist definitive prognoses, and remain sensitive to social change (Luhmann, 2016, p. 746). From this perspective, the frequently reiterated claim in constitutional jurisprudence that legal and constitutional clauses cannot be reduced to mathematical formulas proves at least partially correct. While judicial decisions can be described and anticipated mathematically to a significant degree, such descriptions remain approximations rather than exhaustive representations of legal meaning.

Predictive models, therefore, do not capture law itself, but rather one possible description of its past operations. They translate prior decisions into formal structures that generate probabilistic expectations about future outcomes. The historical evolution of predictive techniques illustrates this limitation clearly: each new generation of models introduces additional variables, corrects earlier assumptions, and adapts to emerging patterns of adjudication. This continuous process of revision underscores the impossibility of fixing legal meaning once and for all, particularly in constitutional contexts where normative texts are continuously reinterpreted in light of changing social expectations.

From a systemic perspective, law must be understood primarily as a mechanism for stabilizing expectations oriented toward the future, rather than as a set of determinate meanings fixed in the present. The function of law is not to eliminate uncertainty, but to render it manageable by producing normative expectations capable of guiding behavior, even though they remain permanently open to frustration (Luhmann, 2016, p. 748). This openness is intensified by the temporalization of normative validity: what is normatively expected today may not be expected tomorrow, as legal meaning evolves in response to social transformations (Luhmann, 2016, p. 751).

This indeterminacy does not imply arbitrariness. On the contrary, the very possibility of prediction demonstrates that legal decision-making exhibits recurrent structures and relatively stable expectations over time. Courts tend to reproduce established patterns for extended periods, making outcomes highly predictable in many cases. Yet such predictability is always probabilistic rather than deterministic. It reflects the provisional stability of social expectations, not the existence of fixed legal meanings immune to change (Luhmann, 2016, p. 752). Predictive models based on artificial intelligence render this structure visible with unprecedented clarity: even the most advanced systems never achieve complete certainty, precisely because legal decision-making remains open to reinterpretation, institutional change, and shifts in social values.

At the same time, the growing use of predictive models produces effects that extend beyond their analytical performance. By transforming past judicial decisions into probabilistic expectations, these systems increasingly inform litigation strategies, case selection, settlement behavior, and institutional planning. In this sense, predictions may acquire a performative dimension, shaping legal practice by influencing perceptions of what outcomes are likely or unlikely. This dynamic entails the risk of reinforcing historically embedded patterns, including structural biases, insofar as models trained on past decisions tend to reproduce the regularities they identify. Predictive systems thus do not merely describe legal behavior but may indirectly participate in its stabilization, raising questions about transparency, contestability, and the governance of algorithmic tools within the legal field.

Seen from this perspective, law cannot be understood as a domain in which uncertainty is progressively eliminated through greater rationalization, whether by idealized judges or by increasingly sophisticated machines. Legal decision-making operates within a temporal structure in which stability and change coexist. Norms do not determine future outcomes; they stabilize expectations that remain permanently exposed to frustration. Predictive models function entirely within this horizon: they may reduce surprise and increase consistency under conditions of relative stability, but they cannot fix legal meaning or exhaust the possibilities of future decisions.

The comparison between Judge Hercules and predictive artificial intelligence thus reveals a shared aspiration to overcome legal uncertainty through rational completeness, an aspiration that proves untenable not because of technical insufficiency, but because of the structural conditions under which law operates. Artificial intelligence does not render law deterministic. Instead, it makes visible a defining feature of legal decision-making: the law's capacity to produce consistency without certainty, order without closure, and expectations without guarantees. What predictive systems ultimately expose is not the end of indeterminacy, but the manner in which uncertainty remains constitutive of law, even where prediction appears most successful.

6. CONCLUSION

The use of artificial intelligence and predictive models in law should not be understood as a technological breakthrough capable of resolving the foundational problems of legal decision-making or of delivering a definitive "correct answer" to legal disputes. Rather, these tools operate as practical instruments that approximate present decisions to patterns observed in past judicial practice. In doing so, they may increase local consistency within the legal system, without transforming it into a fully predictable or deterministic order. Their usefulness lies precisely in the persistence of legal uncertainty: predictive models reduce uncertainty only within the limits imposed by the system's own structure and by the temporal conditions under which legal decisions are produced.

Judge Hercules And Laplace's Demon: Philosophical and Practical Limits of Predictive Artificial Intelligence in Judicial Decision-Making

The comparison developed throughout the article between Judge Hercules and the demon of Laplace clarifies the nature of this limitation. Both figures represent idealized responses to uncertainty grounded in the assumption that complete knowledge of all relevant variables would allow the future to be fully determined. In Laplace's thought experiment, perfect knowledge of the forces and positions governing nature would dissolve uncertainty entirely; in Dworkin's theoretical model, exhaustive knowledge of law would yield a single correct answer to every legal problem. The analysis undertaken here shows that neither aspiration can be fulfilled within the domain of law. Legal decision-making does not operate as a closed system governed by invariant equations, but as a historically situated practice embedded in social, institutional, and temporal dynamics that resist complete formalization.

This insight directly addresses the problem posed: whether contemporary predictive models, supported by artificial intelligence, signal a renewed possibility of overcoming legal indeterminacy and approximating a form of "superhuman" legal rationality. The answer suggested by both theory and empirical research is negative. Although predictive systems may outperform human experts in identifying regularities within large datasets, they do not reconstruct law in its entirety, nor do they eliminate the contingent character of judicial decisions. At most, they offer probabilistic expectations grounded in past patterns, expectations that remain vulnerable to change as legal meanings evolve.

Legal uncertainty, therefore, should not be treated as a defect to be corrected through improved models or greater computational power. It is a consequence of law's detachment from any external, immutable order and of its insertion into a temporally structured society. As social values, institutional practices, and interpretive frameworks change, so too does the meaning attributed to legal norms. Expectations regarding judicial outcomes are necessarily contingent and permanently exposed to frustration. Law operates by projecting normative expectations into an uncertain future, stabilizing behavior not through guarantees, but through what is more or less probable at a given moment.

From this perspective, both the theoretical pursuit of a single correct answer and the practical deployment of predictive artificial intelligence respond to the same structural problem: the impossibility of fully accessing, processing, and fixing all the elements involved in complex decision-making. Their value lies not in transforming law into a perfectly consistent or fully determined system, but in addressing specific functional demands, such as reducing variability, enhancing coordination, or informing strategic choices within bounded contexts.

Artificial intelligence thus remains a tool among others employed by the legal system to manage its own uncertainty. Its growing capacity to identify patterns and regularities may enhance uniformity and reduce surprise, but it does not reveal a hidden determinism beneath legal practice. On the contrary, by demonstrating both the power and the limits of prediction, artificial intelligence ultimately confirms what the analysis of Judge Hercules and the demon of Laplace already suggested: that law continues to operate within an open horizon, where uncertainty is not an obstacle to be overcome, but a constitutive condition of its capacity to adapt, evolve, and respond to social change.

REFERENCES

- 1) Aletras, N., Tsarapatsanis, D., Preotjuc-Pietro, D. and Lamos, V., 2016. Predicting judicial decisions of the European Court of Human Rights: a natural language processing perspective. *PeerJ Computer Science* [online], 2(e93). Available at: <https://doi.org/10.7717/peerj-cs.93>
- 2) Cattani, M., Caldas, I.L., Souza, S.L. and Larosz, K.C., 2016. Deterministic chaos theory: basic concepts. *Revista Brasileira de Ensino de Física*, 39(1).
- 3) Devlin, J., Chang, M.-W., Lee, K. and Toutanova, K., 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *arXiv preprint* [online], arXiv:1810.04805. Available at: <https://arxiv.org/abs/1810.04805>
- 4) Dworkin, R., 2010. *Taking rights seriously*. Cambridge, MA: Harvard University Press.
- 5) Einstein, A., 1936. Physics and reality. *Journal of the Franklin Institute*, 221(3), 349–382. Available at: [https://doi.org/10.1016/S0016-0032\(36\)91047-5](https://doi.org/10.1016/S0016-0032(36)91047-5)
- 6) Frigg, R., Bradley, S., Du, H. and Smith, L., 2014. Laplace's demon and the adventures of his apprentices. *Philosophy of Science* [online], 81(1), 31–59. Available at: <https://doi.org/10.1086/674416>
- 7) Keown, R., 1980. Mathematical models for legal prediction. *Computer/Law Journal*, 2(1), 829–862.
- 8) Kort, F., 1957. Predicting Supreme Court decisions mathematically: a quantitative analysis of the "right to counsel" cases. *American Political Science Review*, 51(1), 1–12.
- 9) Laplace, P.S., 1902. *A philosophical essay on probabilities*. London: Chapman & Hall.
- 10) Lauderdale, B.E. and Clark, T.S., 2012. The Supreme Court's many median justices. *American Political Science Review*, 106(4), 847–866.
- 11) Levy, I., 2022. Hercules on a diet. In: P. Bystranowski and M. Próchnicki, eds., *Judicial decision-making: economic analysis of law in European legal scholarship*. Cham: Springer, 159–177. Available at: https://doi.org/10.1007/978-3-031-11744-2_8

- 12) Li, J., Zhang, Y., Chen, X. and Yu, H., 2019. Research and design on cognitive computing framework for predicting judicial decisions. *Journal of Signal Processing Systems*, 91(10), 1159–1167.
- 13) Loftus, G. and Loftus, E., 1976. *Human memory: the processing of information*. New York: Taylor & Francis.
- 14) Luhmann, N., 2013. *Introduction to systems theory*. Cambridge: Polity Press.
- 15) Luhmann, N., 2016. *Law as a social system*. Oxford: Oxford University Press.
- 16) Luhmann, N., 2018. *Organization and decision*. Cambridge: Cambridge University Press. Available at: <https://doi.org/10.1017/9781108560672>
- 17) Luo, B., Feng, Y., Xu, J., Zhang, X. and Zhao, D., 2017. Learning to predict charges for criminal cases with legal basis. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Stroudsburg, PA: Association for Computational Linguistics, 2727–2736. Available at: <https://arxiv.org/abs/1707.09168>
- 18) Menezes-Neto, E.J. and Clementino, M.B.M., 2022. Using deep learning to predict outcomes of legal appeals better than human experts: a study with data from Brazilian federal courts. *PLOS ONE*, 17(7), e0272287.
- 19) Neves, M., 2016. From legal uncertainty to uncertainty of justice. In: P. Fortes, R. Campos and S. Barbosa, eds., *Contemporary theories of law. Volume I*. Curitiba: Juruá, 43–62.
- 20) Radford, A., Narasimhan, K., Salimans, T. and Sutskever, I., 2018. Improving language understanding by generative pre-training [online]. Available at: <https://www.mikecaptain.com/resources/pdf/GPT-1.pdf>
- 21) Rutkowski, A.-F. and Saunders, C.S., 2019. *Emotional and cognitive overload: the dark side of information technology*. London/New York: Routledge.
- 22) Segal, J.A., 1984. Predicting Supreme Court cases probabilistically: the search and seizure cases, 1962–1981. *American Political Science Review*, 78(4), 891–900.
- 23) Silva, F.G.C.E. and Tizzo, L.G.L., 2023. Habermas' criticism of Dworkin's Judge Hercules and the open society of constitutional interpreters. *Concilium*, 23(4), 263–277.
- 24) Sumathi, S. and Janani, M., 2020. *Neural networks for natural language processing*. Hershey, PA: IGI Global.
- 25) Ulmer, S.S., 1963. Quantitative analysis of judicial processes: some practical and theoretical applications. *Law and Contemporary Problems*, 28, 164–184.
- 26) Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I., 2017. *Attention is all you need*. *Advances in Neural Information Processing Systems*, 30, 5998–6008. Available at: <https://arxiv.org/abs/1706.03762>
- 27) Watson, R.N., 2019. *Cultural evolution and its discontents: cognitive overload, parasitic cultures, and the humanistic cure*. New York: Routledge.
- 28) Wilkins, L., 1973. Information overload: peace or war with the computer. *Journal of Criminal Law and Criminology*, 64(2), 190–197.
- 29) Zhong, H., Guo, Z., Tu, C., Xiao, C., Liu, Z. and Sun, M., 2018. Legal judgment prediction via topological learning. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels: Association for Computational Linguistics, 3540–3549.



There is an Open Access article, distributed under the term of the Creative Commons Attribution – Non Commercial 4.0 International (CC BY-NC 4.0) (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits remixing, adapting and building upon the work for non-commercial use, provided the original work is properly cited.