

PROF. DR. REMO CAPONI
Justice, Italian Supreme Court of Cassation
formerly Professor of Civil Procedure, University of Florence

AI Tools in the Italian Civil Justice System: First Steps and Current Challenges

Japan-Italy Diplomatic Relations 160th Anniversary Commemorative Special Lecture

International Workshop, Ritsumeikan University, Kyoto, 10-11 March 2026

Table of Contents

1. Introduction	2
2. Human Oversight as the Irreducible Core of AI-Assisted Adjudication	3
3. Reframing the Human-AI Relationship.....	7
4. AI as Powerful Extension of Human Cognitive Capacities.....	9
5. The Pioneering Legacy of the Italian Supreme Court’s CED	10
6. From Italgire to the Digital Age: Structural Evolution	12
7. The Contemporary Challenge: Caseloads, Efficiency, and the Turn to AI	13
8. The Epistemological Challenge: Distinguishing Fact from Law	15
9. The European Regulatory Framework: The EU AI Act	16
10. Italy’s National AI Legislation: Law No. 132/2025	18
11. The Recommendations of the High Council of the Judiciary	20
12. Judicial Training in the Age of AI.....	22
13. From Theory to Practice: National Rollout of Microsoft Copilot 365	23
14. AI in Action: PRODIGIT versus CED—Technical Challenges	26
15. Pilot Projects Across the Italian Judiciary.....	28
16. Emerging Case Law: Italian Courts Confront AI	30
17. Comparative Remarks: Four Structural Tensions	32
(a) Executive Influence versus Judicial Independence.....	32
(b) Proprietary Platform Centralisation versus Digital Sovereignty	33
(c) The Multi-Level Regulatory Overload	33
(d) Epistemological Transparency versus Adjudicative Autonomy.....	34
18. Closing Remarks	35

1. Introduction

I am very grateful to Professor Masahisa Deguchi and to the Konrad-Adenauer-Stiftung (Rule of Law Programme Asia) for the kind invitation to address this distinguished assembly. The present workshop builds upon an admirable tradition of international scholarly exchange hosted at Ritsumeikan University — from the 2017 symposium on the protection of fundamental rights in European and Asian law, through the 2019 conference on Europe and Asia as a legal area for peace and the rule of law,¹ to the 2023 and 2024 workshops on e-justice, information technology, and consumer protection. It is a privilege to contribute to this ongoing dialogue between legal cultures.

This lecture is delivered in the year that marks the 160th anniversary of the establishment of formal diplomatic relations between Italy and Japan. On 25 August 1866, Captain Vittorio Arminjon signed the Treaty of Friendship and Commerce in Edo, inaugurating a bilateral relationship that has since developed into what both governments now describe as a special strategic partnership.²

In January 2026, Prime Minister Takaichi Sanae and President of the Council of Ministers Giorgia Meloni reaffirmed their shared commitment to cooperation in advanced technology fields, including artificial intelligence.³ The present paper aspires to make a modest scholarly contribution to this spirit of bilateral intellectual exchange.

The paper was originally announced under the programme title “The Italian Supreme Court: Structural and Procedural Challenges in the Age of AI.” The scope has been broadened to encompass the Italian civil justice system as a whole, for a reason that is itself instructive: artificial intelligence does not respect the organisational boundaries of the judiciary. Its introduction affects courts of first instance, appellate courts, and supreme courts alike; it touches upon judicial training, regulatory governance, digital infrastructure, and the very epistemology of adjudication. The revised title reflects this wider aperture, while keeping the focus primarily on the Supreme Court (*Corte di cassazione*).

¹ M. Deguchi and K. Yakushiji (eds.), *Europe and Asia as a Legal Area for Fundamental Rights*, Springer 2023.

² On the history of Italy-Japan relations, see the comprehensive account on the website of the Ministry of Foreign Affairs of Japan, *Italy–Japan Relations (Archives)*, available at <https://www.mofa.go.jp/region/europe/italy/archives.html>.

³ Japan-Italy Joint Statement, 16 January 2026, available at <https://japan.kantei.go.jp/104/diplomatic/202601/16italy.html>.

I was a professor of civil procedure, and I currently serve as a justice of the Italian Supreme Court of Cassation. This dual perspective — academic and judicial — informs the paper’s approach: it seeks to combine doctrinal analysis with first-hand institutional knowledge, placing the Italian experience within the broader European regulatory framework and in dialogue with the digital justice systems of the jurisdictions represented at this workshop, particularly Switzerland, South Korea, Taiwan, the People’s Republic of China, and the United States.

The paper proceeds as follows. It begins with a theoretical foundation, articulating the principle of human oversight—the Italian *riserva di umanità*—as the irreducible normative core of AI-assisted adjudication, while arguing against a purely defensive posture toward technological innovation. It then adopts a historical and genealogical perspective centred on the Supreme Court: the emergence of the Court’s Electronic Documentation Centre and the Italgire infrastructure, the structural pressures of caseloads and efficiency, and the epistemological constraint posed by the fact–law distinction. Against this institutional backdrop, the paper maps the multi-level regulatory landscape: the EU AI Act, the Italian national legislation, and the institutional responses of the High Council of the Judiciary and the Higher School of the Judiciary. The analysis then turns to current implementation, focusing on the nationwide deployment of Microsoft Copilot 365. Concrete AI applications are subsequently examined: the PRODIGIT project, pilot projects across Italian courts, and the emerging case law. The paper concludes by identifying four structural tensions that crystallise the Italian experience and may serve as analytical categories for comparative reflection: the balance between executive governance and judicial independence; digital sovereignty versus vendor dependence; the risk of multi-level regulatory overload; and the epistemological challenge that AI poses to the fact/law distinction on which the architecture of appellate review depends.

2. Human Oversight as the Irreducible Core of AI-Assisted Adjudication

The starting point for any serious analysis of artificial intelligence in the judiciary must be the principle of human oversight — known in Italian legal doctrine as the *riserva di umanità*, literally the “reservation of humanity.” This principle holds that certain decisions, above all judicial ones, must remain irreducibly reserved to human beings and cannot be delegated to AI systems, however sophisticated.

The concept designates the minimum, irreducible sphere of competence that must remain exclusively within human hands in the exercise of public

functions, particularly judicial decision-making. As it has been put, the *riserva di umanità* defines the minimum sphere of exclusive prerogative of the human being in the exercise of state functions.⁴ In the judicial sphere, this reservation acquires reinforced intensity: the Constitution demands not merely the humanity of the decision but the humanity of the decider — the very personhood of the judge, to whom the attributes of impartiality, independence, and responsibility are constitutionally ascribed.

The intellectual origins of this principle lie at the intersection of three constitutional traditions. First, human dignity as a foundational value — Article 2 of the Italian Constitution (“The Republic recognises and guarantees the inviolable rights of the person”), Article 1 of the German *Grundgesetz* (“Human dignity shall be inviolable”), and Article 1 of the EU Charter of Fundamental Rights. Dignity operates here as a barrier against reducing individuals to objects of automated state processes. Second, due process and the right to a fair trial: Article 6 of the European Convention on Human Rights and Article 47 of the EU Charter require adjudication by an “independent and impartial tribunal established by law,” presupposing a human adjudicator capable of giving reasons, exercising judgment, and bearing personal responsibility. Third, judicial independence: Articles 101 and 102 of the Italian Constitution vest judicial power exclusively in human judges, who are “subject only to the law” — a formulation that implicitly excludes subordination to algorithmic outputs, which would amount to an improper external determination of judicial reasoning by a machine.

No single English term captures the full normative content of the Italian *riserva di umanità*. The concept maps onto a constellation of related but distinct formulations. Human oversight, employed by Article 14 of the EU AI Act, represents the broadest regulatory expression. The EU Ethics Guidelines for Trustworthy AI (2019) operationalise this through three graduated models: Human-in-the-Loop (HITL), where humans actively participate in each decision cycle; Human-on-the-Loop (HOTL), where humans supervise and intervene when necessary; and Human-in-Command (HIC), where humans retain ultimate, non-delegable decisional authority.⁵ Meaningful human control, originating in the 2013 UK policy discourse on autonomous weapon systems, has since migrated into civilian AI governance.⁶ The GDPR’s right not to be subject to a decision based

⁴ G. Gallone, *Riserva di umanità e funzioni amministrative*, Cedam 2023.

⁵ European Commission, High-Level Expert Group on AI, *Ethics Guidelines for Trustworthy AI*, 2019.

⁶ F. Santoni de Sio and J. van den Hoven, “Meaningful Human Control over Autonomous Systems: A Philosophical Account,” *Frontiers in Robotics and AI*, 2018.

solely on automated processing (Article 22) represents the data-protection dimension. The CEPEJ Charter's Principle 5, under user control, adds the sector-specific judicial dimension.⁷

These overlapping formulations converge on a common insight that Álvaro Pérez Ragone, writing from the Latin American perspective, articulates through the analytical lens of Niklas Luhmann's distinction between *trust* and *confidence*. Confidence is a passive, systemic expectation: citizens assume the institutional machinery will function without actively engaging with it. Trust, by contrast, is relational and active: it involves a willingness to accept vulnerability by submitting one's dispute to human adjudicators whose integrity and impartiality one relies upon. Pérez Ragone argues that algorithmic opacity, the "black box" problem, erodes both dimensions simultaneously: it undermines system-level confidence in the institution and destroys the inter-personal trust between litigant and adjudicator.⁸

Pérez Ragone's theoretical framework finds concrete illustration in the Latin American landscape he surveys. Brazil's judiciary has deployed some 140 AI tools under the *Justice 4.0* programme launched in 2020 by the National Council of Justice (*Conselho Nacional de Justiça*), encompassing case classification, similarity detection, mass-litigation identification, and workload forecasting — a scale of adoption unmatched in any other jurisdiction.⁹ In Colombia, the Constitutional Court's PretorIA system — developed in collaboration with the University of Buenos Aires's IALAB and operational since 2020 — processes approximately 4,500 *tutela* (constitutional protection) filings daily, assisting the Court in docket selection by identifying cases of particular urgency or systemic importance; the adjudication itself remains entirely in human hands.¹⁰ In

⁷ CEPEJ, *European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and their Environment*, December 2018, Principle 5.

⁸ Á. Pérez Ragone, *Supreme Courts with a Human Face: Citizen Trust and Confidence in Latin America Amidst Digitalization and AI*, presentation delivered at the Kyoto International Workshop. Drawing on N. Luhmann, *Trust and Power*, Chichester, 1979.

⁹ On the Brazilian experience, see the National CNJ Survey of AI Models and the Oxford Institute of Technology and Justice, *Country Profile: Brazil (2025)*, available at <https://www.techandjustice.bsg.ox.ac.uk/research/brazil>. The 140 figure refers to AI projects catalogued in the CNJ survey; the number of fully operational systems is smaller.

¹⁰ On PretorIA, see Oxford Institute of Technology and Justice, *Country Profile: Colombia (2025)*, available at <https://www.techandjustice.bsg.ox.ac.uk/research/colombia>; P. Rueda Saiz, *Docket Selection and Judicial Responsiveness: The Use of AI in the Colombian Constitutional Court*, 30 *Wm. & Mary Bill Rts. J.* 371 (2021). In Ruling T-323 of 2024, the Colombian Constitutional Court addressed for the first time the use of generative AI (ChatGPT 3.5) by a lower-court judge in *tutela* proceedings, establishing guiding principles centred on transparency, non-substitution of human rationality, and accountability.

Mexico, initial experimentation with generative AI tools for preliminary legal research is reported, though still at an early stage. Each of these experiences illustrates, in a context of structurally fragile judicial legitimacy the human-in-the-loop safeguard is not a luxury of consolidated democracies but a universal institutional necessity.¹¹

What the Italian *riserva di umanità* adds to all of these is a distinctive positive character. It does not merely prohibit full automation but affirms the constitutional necessity of the human judge as person — a being endowed with irreplaceable evaluative, hermeneutic, and empathic capacities. As Giovanni Maria Flick, President Emeritus of the Italian Constitutional Court, has argued, the price of humanity includes emotivity and empathy, reasonable doubt, a knowledge that goes beyond the appearance of algorithmic knowledge.¹²

The *riserva di umanità* serves interconnected functions. It safeguards judicial independence and impartiality: a judge who uncritically follows AI-generated recommendations ceases to exercise independent judgment. It protects the right to a fair trial and due process: Article 6 ECHR presupposes a human adjudicator capable of hearing, understanding, and responding to individual circumstances. It ensures accountability and personal responsibility: Article 28 of the Italian Constitution presupposes human agents capable of bearing personal responsibility for their decisions. And it maintains the democratic legitimacy of adjudication: justice administered “in the name of the people” (Article 101 Italian Constitution) requires a human intermediary who can translate democratic legal norms into individualised determinations.

This principle is of course not uniquely Italian. It resonates across the legal orders represented at this workshop. In Switzerland, the Federal Supreme Court has adopted what Judge Christoph Hurni calls the “Iron Law”: AI is a tool, not a judge.¹³ In South Korea, the Framework Act on the Development of AI and Establishment of Trust (the “AI Basic Act,” enacted January 2025, effective January 2026) mandates human

¹¹ Á. Pérez Ragone, *Supreme Courts with a Human Face*, supra. Pérez Ragone shows in the end that in Latin America, where judicial legitimacy suffers from a structural deficit rooted in authoritarian legacies, the digitalization of justice carries heightened risks: algorithmic opacity may compound, rather than cure, the deep-seated public distrust of judicial institutions. His conceptual framework distinguishes input legitimacy (control over data integrity), throughput legitimacy (understanding of algorithmic procedures), and output legitimacy (accuracy of results compared with human adjudication).

¹² G.M. Flick, C. Flick, *L'algoritmo d'oro e la torre di Babele: Il mito dell'informatica*, Baldini Castoldi 2022.

¹³ Ch. Hurni, “Digital Justice in Switzerland and the Use of AI at the Swiss Federal Supreme Court,” presentation delivered at the Kyoto International Workshop.

oversight mechanisms and fundamental rights impact assessments; the General Act on Public Administration (2021) explicitly excludes discretionary decisions from full automation.¹⁴ In Taiwan, the Judicial Yuan's AI-assisted judgment drafting system explicitly generates "drafts" for judicial reference, with judges retaining full authority.¹⁵ In the People's Republic of China, the Supreme People's Court's 2022 Opinions articulate a Principle of Supporting Adjudication: "AI shall not make judicial decisions substituting for the judge in any case."¹⁶ In the United States, the *State v. Loomis* framework acknowledges AI risk-assessment tools while preserving the sentencing judge's ultimate discretion.¹⁷ The core insight is universally shared: there exists an irreducible domain of human judgment that no AI system, however sophisticated, can legitimately occupy in the administration of justice.

3. Reframing the Human-AI Relationship

Having established the normative foundation, it is essential to resist the temptation of a purely defensive stance. I argue that an exclusively cautionary posture toward artificial intelligence risks misunderstanding both the nature of technology and the nature of the human subject it purports to protect.

The starting point is an anthropological observation. The human being is not a stable, self-contained entity confronting technology from a position of sovereign autonomy.

Friedrich Nietzsche described the human as *das noch nicht festgestellte Thier*, "the animal not yet defined" (*Beyond Good and Evil*, § 62). Arnold Gehlen, drawing on Johann Gottfried Herder's earlier anthropological insight, elaborated this intuition into the systematic concept of the *Mängelwesen*, the "deficient being": a creature whose biological incompleteness is compensated through culture, institutions, and technique.¹⁸ The human

¹⁴ Republic of Korea, Framework Act on the Development of AI and Establishment of Trust, enacted January 2025; General Act on Public Administration, 2021. See also Korean judiciary's *Guidelines on AI Use in the Judiciary*, February 2025.

¹⁵ Taiwan Judicial Yuan, press release of 27 August 2023 on the AI-assisted judgment drafting system.

¹⁶ Supreme People's Court of the People's Republic of China, *Opinions on Regulating and Strengthening the Applications of AI in Judicial Fields*, December 2022.

¹⁷ *State v. Loomis*, 881 N.W.2d 749 (Wis. 2016).

¹⁸ F. Nietzsche, *Beyond Good and Evil: Prelude to a Philosophy of the Future*, § 62, Judith Norman tr, Rolf-Peter Horstmann and Judith Norman eds, Cambridge University Press 2002; A. Gehlen, *Man: His Nature and Place in the World*, Clare McMillan and Karl Pillemer trs, Columbia University Press 1988 (original: *Der Mensch. Seine Natur und seine*

being's very identity is thus constantly shaped and reshaped by the tools it creates.

One may recall the celebrated opening sequence of Stanley Kubrick's *2001: A Space Odyssey*, "The Dawn of Man," in which an ape's first use of a bone as a tool marks a threshold in hominisation. From this perspective, we are less stable human "beings" than human "becomings" — creatures continually shaped and re-shaped by our technological milieu.¹⁹

Artificial intelligence represents a historically pivotal moment in this ongoing process of co-constitution — a juncture comparable, in its transformative potential, to the invention of the Greek alphabet and the Indo-Arabic numeral system. Binary code is the latest frontier of what one might call the "extraordinary machine of language" — the apparatus through which human thought has progressively objectified itself, from pictograms to phonetic scripts, from the printing press to the digital revolution, and now to systems capable of manipulating linguistic symbols with a fluency that simulates understanding.

The philosophical genealogy of this insight has been developed with particular rigour by one of the most prominent Italian philosophers, Carlo Sini, whose work on writing and technology constitutes a sustained meditation on the constitutive relationship between symbolic practices and the formation of the human subject. In his seminal *Etica della scrittura* (1992),²⁰ Sini demonstrated that alphabetic writing is not a neutral tool for recording pre-existing thought but the very condition that produced the logical mind of Western civilisation — what he calls "the content of form": the technological form of the alphabet does not merely shape but "literally" (*pour cause!*) generates the rational subject who employs it. Each great writing practice — from pictograms to phonetic scripts, from printing to mathematical notation — does not merely transmit knowledge but reshapes the cognitive architecture of the human being who inhabits it.

Stellung in der Welt, 1940). Gehlen systematised an insight that traces to J.G. Herder, *Abhandlung über den Ursprung der Sprache*, 1772.

¹⁹ S. Kubrick (dir.), *2001: A Space Odyssey*, Metro-Goldwyn-Mayer, 1968, opening sequence ("The Dawn of Man"). Also on Youtube.com. Keywords: "2001: A Space Odyssey - The Dawn of Man."

²⁰ C. Sini, *Ethics of Writing*, Silvia Benso with Brian Schroeder trs, State University of New York Press, 2009. Original version: C. Sini, *Etica della scrittura*, Milan, 1992 (reprinted Mimesis, 2009). Sini's genealogy of the "logical mind" through alphabetic practice converges with the anthropological premises of Gehlen and Nietzsche while adding a distinctively semiotic and pragmatist dimension drawn from Peirce.

Plato's *Phaedrus* offers a salutary reminder. In the celebrated passage, Socrates recounts the Egyptian myth of Theuth, inventor of writing, whose gift was rejected by King Thamus on the ground that it would produce "forgetfulness in the learners' souls" and merely the "appearance of wisdom, not true wisdom."²¹ The structural homology with contemporary anxieties about AI is striking: the fear that a new technology will displace authentic human capacities — memory, judgment, understanding — and substitute a mere simulacrum. History has repeatedly vindicated the adopters over the sceptics.

This observation generates a fundamental philosophical puzzle: how can one logically demand that control remain "human" when the human subject itself is constantly being redefined by the very technology it purports to control? The circle appears vicious. Yet it is precisely within this circle that the normative work must be done. Rather than positing a fixed human "essence" that technology threatens, one may recognise that each great technological revolution — writing, printing, computing — has both expanded and reconfigured the human.

4. AI as Powerful Extension of Human Cognitive Capacities

Artificial intelligence will inevitably retrace the contours of what it means to be a human judge, just as digitalisation has already transformed judicial work and the printing press transformed legal science.

The critical insight is that AI is not an autonomous alien entity. It is objectified human labour — the crystallised product of human ingenuity, human data, human design choices, and human values (or their absence). To regulate AI is, in this sense, to regulate the memory and accumulated work of human civilisation as it has been encoded in computational form. The machines are never truly "other"; they are the sediment of ourselves.

This recognition does not dissolve the need for the *riserva di umanità* — quite the contrary. It recasts it in more precise terms. The crucial distinction is not between "human" and "machine" as ontological categories, but between *calculation* (*Berechnung*) and *understanding* (*Verstehen*): between the syntactic manipulation of symbols and the semantic grasp of meaning.²² Current AI systems — including the most powerful large language models — operate on the plane of statistical correlation and pattern

²¹ Plato, *Phaedrus*, 275a–b.

²² The distinction has deep roots in the continental hermeneutic tradition. For its application to the AI context, see R. Caponi, "Oralità e scrittura del diritto, intelligenza artificiale," in *Rivista trimestrale di diritto e procedura civile*, 2024, 367 ff.

recognition. They do not *understand* in any philosophically robust sense; they *compute*. A sentence produced by a language model may be syntactically impeccable and pragmatically useful, yet it lacks the intentional relationship to reality that characterises genuine comprehension.

Human oversight, then, is the site of responsibility — the juncture where calculation meets judgment, where statistical probability meets normative evaluation, where the realm of necessity encounters the realm of possibility. Machines operate in the domain of necessity and computation; legal and ethical decisions inhabit the domain of possibility and accountability for the future. The judge who decides a case does not merely compute an output from a set of inputs: the judge evaluates, interprets, balances competing values, exercises practical wisdom (*phronesis*), and assumes personal responsibility for the outcome. It is this irreducible dimension of responsibility — not a metaphysical claim about human nature — that grounds the *riserva di umanità*.

Two remarks should be kept constantly in mind before turning to institutional settings, norms, and compliance. First, the human element that the *riserva di umanità* seeks to preserve is not an ahistorical remainder opposed to technique. Judicial reasoning has always been enacted through material and symbolic supports—writing practices, archives, citation systems, and institutional routines—which shape what counts as a reason and how reasons are stabilised. Second, this is why the influence of language technologies is, to a significant extent, inevitable: once a tool becomes a routine medium of work, it quietly reconfigures cognitive habits and argumentative pathways. The normative task is therefore to make such influence visible, contestable, and governable.

From this standpoint, a brief historical and genealogical sketch is not a digression but a prerequisite. To understand what AI can and cannot responsibly do in adjudication, one must first grasp the documentary infrastructures and organisational constraints through which judicial work—on this occasion at the Italian Supreme Court—has long been mediated.

5. The Pioneering Legacy of the Italian Supreme Court's CED

That broader perspective can be conveyed through a single, emblematic episode: the creation of the Electronic Documentation Centre (CED), conceived in the late 1960s and institutionalised in the early 1970s in Rome. The CED is not merely an IT department. It is the institutional technology through which the Court's jurisprudence is collected, indexed, retrieved, and—crucially—made writable and governable as a body of reasons. Seen genealogically, the CED/Italgjure trajectory shows how 'judicial knowledge' has been progressively externalised into documentary

infrastructures, thereby reshaping both the possibilities and the limits of adjudication. Against this background, the story of AI in the Italian judiciary does not begin in 2025; it is the latest chapter in an institutional trajectory spanning over half a century. Over time, the CED helped transform the Italian judiciary from a paper-bound bureaucracy into a pioneer of legal informatics, and it continues to frame the conditions under which contemporary AI tools are introduced.

The CED's history is inseparable from the Court's structural workload. The Italian Court of Cassation—rooted in the post-Revolutionary French model—was designed to secure the uniform observance of the law in a unified nation. Moreover, Article 111 of the Constitution guarantees a right of appeal to the Court of Cassation for violations of law. Together, these institutional features generated a sustained, high-volume inflow of cases, making robust documentary infrastructures necessary well before the advent of contemporary computing.

The driving force behind the CED's pioneering role was Renato Borruso, a judge whose vision went beyond the mere mechanisation of archives. Borruso and his colleagues treated access to precedent as a condition of legal legitimacy: principles had to be retrievable, contestable, and consistently applied. In the Italian context, this required the systematic digitisation of the *massime*—abstracts that distil the core holding (*ratio decidendi*) of a judgment. In February 1973, the Court publicly inaugurated the Italgire system, establishing an institutional repository that integrated case law, legislation, and, over time, selected doctrinal materials.²³

To situate the CED historically, it is useful to recall the broader landscape of early legal informatics. In the United States, experimental full-text retrieval projects for legal materials began in the 1960s; from 1973, Lexis helped inaugurate the commercial model that would later dominate the Anglo-American market. Parallel initiatives emerged in other European jurisdictions. Italgire was distinctive for its institutional integration: it was not a commercial service, but a public-sector infrastructure embedded within the Supreme Court to serve its constitutional function of *nomofilachia*—the uniform interpretation of the law. Its dual focus—covering both case law (*massime*) and legislation, and later also selected scholarly commentary—set it apart from many contemporaneous systems.

The fact that judges—rather than technologists or commercial actors—stood at the helm of the CED is a distinctive feature of the Italian

²³ On the inauguration of Italgire in 1973, see the account in the ITTIG Newsletter, 1973; and G. Oberto, "Appunti per un corso di informatica giuridica, available at www.giacomooberto.com.

trajectory in legal informatics. This institutional choice reflects a broader governance claim: the organisation of legal knowledge is too important to be left entirely to the market.

6. From Italgire to the Digital Age: Structural Evolution

The normative framework governing the CED has evolved through successive legislative interventions establishing the CED's organisational structure and transferring certain functions to the Ministry of Justice. A significant bifurcation emerged between legal informatics (*informatica giuridica*) — concerned with the substance of legal documentation — and judicial informatics (*informatica giudiziaria*) — concerned with the technological infrastructure of the courts. The CED straddles both domains.

Today, Italgire contains approximately 35 million documents, spanning case law, legislation, selected scholarly commentary, and administrative materials. In current judicial AI deployments, it has been identified as the certified institutional repository for legal retrieval—an important choice, as it anchors AI-assisted research in a quality-controlled source rather than the open internet, with its attendant risks of error and confabulation (hallucinations).

The CED is also engaged in AI-related research and experimentation through institutional collaborations with academic partners, aimed at organising digital legal materials and developing tools for the extraction and representation of legal knowledge. This continuity—linking half a century of documentary infrastructure to present AI initiatives—matters for governance: it suggests that the most accountable forms of AI assistance are those anchored in public, court-embedded repositories and in research arrangements that preserve auditability and institutional responsibility.²⁴

Yet documentary infrastructure has always been built under a second, more prosaic constraint: volume. The CED was, from the outset, also a response to the Court's chronic inflow of appeals and to the need to stabilise a coherent jurisprudence under quantitative pressure. That same pressure—now measured in tens of thousands of pending cases—forms the immediate background against which contemporary AI tools are being considered.

²⁴ Giuseppe Atzori, 'Prove di giustizia predittiva in Italia', 7(2) *Journal of Ethics and Legal Technologies*, 2025, § 3.5, reporting the strategic collaboration agreement signed in 2021 between the Court of Cassation, through the CED, and the IUSS Pavia, aimed at organising digital jurisprudential and legislative materials through legal analytics and AI, with an initial five-year term, renewable for a further five years.

7. The Contemporary Challenge: Caseloads, Efficiency, and the Turn to AI

The structural overload of the Italian Supreme Court of Cassation is, to a significant extent, a consequence of the constitutional architecture. Article 111(7) of the Constitution provides that, against judicial judgments—and against measures affecting personal liberty—an appeal to the Court of Cassation for violations of law is always available. Unlike the certiorari model of the United States Supreme Court, or the leave-to-appeal requirements of many common-law systems, this design tends to generate a court of last resort with a structurally high inflow.

The statistical data from the latest Report on the Administration of Justice (30 January 2026) confirm a significant improvement but also underscore the scale of the challenge. As of 31 December 2025, the Court has 79,812 civil cases pending — down from 87,380 at the end of 2024 and from the peak of approximately 120,000 at the end of 2020. New filings in 2025 totalled 26,345, while 34,062 cases were resolved, yielding a *clearance rate* (or turnover index) of 129.3%.²⁵

The *disposition time* — the CEPEJ indicator measuring the predicted time to clear the backlog assuming constant case flows — has fallen from the 2019 baseline of 1,302 days to 863 days in 2025, representing a reduction of 33.7%. The PNRR (National Recovery and Resilience Plan) target of 977 days, set for June 2026, has been surpassed with eighteen months to spare.²⁶ The average actual duration of completed civil proceedings, however, remains at 1,191 days (approximately three years and three months), though this figure includes the resolution of older legacy cases.

These improvements reflect several factors: the establishment of the *Ufficio per il Processo* (Office for the Process)—a clerking support structure assisting judges; the reorganisation of sectional screening offices; the expansion of electronic filing and the publication of decisions by presiding judges; and, not least, the sustained human effort of judges who routinely

²⁵ Report on the Administration of Justice, delivered by the First President of the Court of Cassation, Pasquale D’Ascola, 30 January 2026 (*Relazione sulla amministrazione della giustizia nell’anno 2025*) available at cortedicassazione.it

²⁶ The PNRR (*Piano Nazionale di Ripresa e Resilienza*) is Italy’s National Recovery and Resilience Plan: the domestic programme that allocates EU Next Generation EU Recovery and Resilience Facility funds to reforms and public investments (including justice and digitalisation). The facility was adopted in the aftermath of the pandemic. The PNRR set a baseline of 1,306 days (2019) and a target of 977 days by June 2026 for the disposition time of civil proceedings in the Court of Cassation.

handle between 200 and 260 cases per year—a workload that entails significant physical and psychological costs.

The Tax Division (*Sezione Tributaria*) remains a major driver of the backlog: in the First President’s address, tax cases still represent around 40% of the Court’s pending docket. More generally, tax-settlement schemes and transitional measures introduced in recent years, including those connected with Law No. 130 of 31 August 2022 and Law No. 197 of 29 December 2022—may affect the observable composition and dynamics of inflows and outflows, complicating trend analysis.

Despite these improvements, the fundamental structural problem remains unresolved: the Court’s access remains effectively unrestricted. The General Assemblies of the Court in 1999, 2015, and 2025 have debated reforms to the appellate filter, but no consensus for restriction has emerged²⁷.

²⁷ For an in-depth analysis see R. Caponi, “Nomofilachia under Strain: Structural and Procedural Challenges in the Italian Supreme Court,” in *ZJP International* 29 (2024), 2025, 39 ff. The most recent measures, enacted under the 2021–2022 civil justice reform, introduced an accelerated dismissal procedure (Article 380-bis CCP), applicable in all proceedings before the civil sections. Under this mechanism, a preliminary screening identifies cases likely to be dismissed due to inadmissibility, procedural bar, or a manifest lack of merit. For such cases, a single judge (the section president or a delegate) issues a draft opinion proposing dismissal or rejection. If the appellant does not request referral to the judicial panel under the chamber proceedings within a brief time limit, the appeal is deemed tacitly withdrawn and proceedings are terminated. The incentive for compliance lies in avoiding a monetary sanction equivalent to the initial filing fee (a sanction that would otherwise be imposed if the appeal were dismissed or rejected by the panel). Conversely, if the appellant requests a decision and the judicial panel confirms the draft opinion of the delegate judge, severe monetary sanctions for abuse of process are imposed. This mechanism has provoked immediate scholarly criticism, particularly with regard to constitutionality: doubts have been raised as to its compatibility with impartiality, given that the judge who drafts the preliminary opinion may sit on the judicial panel, and with the right to defence, due to the strong deterrent effect of financial sanctions. The most trenchant critique concerns the “monocratic drift”—that is, the shift towards decision-making by individual judges rather than collegial bodies—within an institution historically rooted in collegiality, and its transformation from a judicial body into a manager of judicial resources. It is premature to provide a conclusive assessment of the accelerated dismissal procedure. It is evident that the mechanism has thus far not demonstrated a great degree of effectiveness. The outcome can be attributed, in large part, to the circumspect approach adopted during its initial implementation, which was aimed at avoiding widespread opposition from the legal profession. Historically, the Italian legislator has been highly sensitive to such opposition, reflecting its susceptibility to political pressures from the legal profession. This opposition is evidenced by the criticism from civil procedure scholars, who argue that Article 380-bis CCP poses an increased risk to the integrity of procedural guarantees. Conversely, for a political system genuinely committed to enhancing civil justice and capable of counterbalancing the corporatist interests of the legal profession, the experience with the simplified procedure under Article 380-bis CCP could serve as a model for designing a case selection mechanism,

It is against this background that the turn to AI must be understood. For the Kyoto audience: the Italian Supreme Court's interest in artificial intelligence is not a matter of technological fashion. It is a structural response to a constitutional challenge, the challenge of maintaining the quality of justice under conditions of relentless quantitative pressure.

8. The Epistemological Challenge: Distinguishing Fact from Law

Yet the Court's workload is not only quantitative. It is also shaped by the boundaries of cassation review, where the line between fact and law is constantly contested.

The distinction between questions of fact and questions of law constitutes a perennial challenge for any supreme court operating within a legal system that confines its competence to errors of law. In Italy, Article 360(1)(5) of the Code of Civil Procedure was rewritten in 2012—restricting review of the reasoning on facts to the “omitted examination of a decisive fact”—and this intensified the difficulty, as the landmark rulings of the Supreme Court Joint Divisions (*Sezioni Unite*), Nos. 8053 and 8054 of 2014, made clear.

The continental hermeneutic tradition and the analytical philosophy of evidence converge on a fundamental insight: the distinction between fact and law cannot be drawn with mathematical precision. It ultimately reflects a choice of judicial policy—namely, how decisional authority should be allocated between the court below and the court of review.²⁸

AI may offer a promising contribution here. Natural-language processing tools can assist in analysing the reasoning of lower-court decisions, identifying motivations that are “merely apparent” (*motivazione meramente apparente*), internally contradictory, or otherwise defective in ways that, under the post-2012 framework, remain reviewable as legal errors. AI-enabled cognitive filtering could support the Court's screening function, helping to distinguish cases that raise genuine issues of legal principle from those that are, in substance, disguised requests for factual re-examination. This potential aligns with what one may call a “particularist” paradigm at the level of the Supreme Court—a model in which the Court

which would require amending Article 111 of the Constitution. The initial step in this direction should be the conversion of the preliminary opinion for dismissal (or rejection) into a definitive decision adopted by the panel, given that at present such an opinion results in the termination of cassation proceedings only if accepted by the appellant.

²⁸ J.A. Jolowicz, *On Civil Procedure*, Cambridge, 2000; R.J. Allen and M.S. Pardo, “The Myth of the Law-Fact Distinction,” in *Northwestern University Law Review*, Vol. 97, 2003.

attends closely to the specific contours of each case, rather than processing appeals mechanically.²⁹

9. The European Regulatory Framework: The EU AI Act

To move from a structurally motivated interest in AI to concrete deployment, the Italian Supreme Court's trajectory must now be situated within a regulatory framework. Within the European Union, that lens is provided, first of all, by the AI Act.

Regulation (EU) 2024/1689 of 13 June 2024, commonly known as the AI Act, is widely described as the world's first comprehensive legislative framework for the regulation of artificial intelligence. For an audience primarily situated outside the European Union, its main features merit a concise but precise exposition.

The AI Act adopts a risk-based approach, classifying AI systems into four tiers of regulatory intensity. Systems posing unacceptable risk—including social scoring by public authorities and certain forms of real-time remote biometric identification—are categorically prohibited (Article 5). High-risk systems are subject to stringent requirements concerning data quality, documentation, transparency, human oversight, accuracy, robustness, and cybersecurity (Articles 8–15). Certain other systems are subject to targeted transparency obligations (Article 50). Minimal- or no-risk systems remain largely unregulated.

For the judiciary, the decisive provision is Annex III, point 8(a), which classifies as high-risk any AI system “intended to be used by a judicial authority or on their behalf to assist a judicial authority in researching and interpreting facts and the law and in applying the law to a concrete set of facts, or to be used in a similar way in alternative dispute resolution.” Recital 61 justifies this with the “potential significant impact on democracy, the rule of law, individual freedoms, and the right to an effective remedy and to a fair trial,” while also clarifying that the high-risk classification should not extend to purely ancillary administrative tasks that do not affect adjudication in individual cases.³⁰

The high-risk classification entails that judicial AI systems must comply with the full panoply of requirements under Title III, Chapter 2—including

²⁹ M. Taruffo, “Un vertice giudiziario astratto,” in *Il Anuario de la Facultad de Derecho de la Universidad Autónoma de Madrid*, no. 22, 2018; R. Caponi, “Nomofilachia under Strain,” *supra*.

³⁰ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (AI Act), Recital 61.

risk management (Article 9), data governance (Article 10), technical documentation (Article 11), record-keeping (Article 12), transparency and provision of information to deployers (Article 13), human oversight (Article 14), and accuracy, robustness, and cybersecurity (Article 15). Article 14 on human oversight is of particular importance: it requires that high-risk AI systems be designed so that natural persons can effectively oversee their functioning, duly monitor their operation, correctly interpret outputs, remain aware of automation bias, decide not to use the system (or to disregard, override, or reverse its output), and, where necessary, intervene or interrupt the system through a “stop” button or an equivalent procedure.

A significant qualification arises from Article 6(3), which allows an AI system listed in Annex III not to be treated as high-risk where it does not pose a significant risk of harm to health, safety, or fundamental rights, including by not materially influencing the outcome of decision-making. Among the conditions considered relevant are whether the system performs a “narrow procedural task” or is intended to improve the result of a previously completed human activity. This derogation has generated debate regarding its scope: does it cover purely administrative functions (e.g., case management, scheduling, anonymisation), or might it be invoked to exclude certain substantive support tools? The Commission’s guidelines on the practical implementation of Article 6 (including illustrative examples of use cases that are high-risk or not high-risk) are due no later than 2 February 2026 and will be closely scrutinised.³¹

Article 5 also establishes absolute prohibitions of direct relevance to judicial settings, especially in criminal-justice or security-adjacent contexts. These include AI systems that make risk assessments of natural persons in order to assess or predict their risk of committing a criminal offence solely on the basis of profiling or of an assessment of their personality traits and characteristics (Article 5(1)(d)), as well as the building or expansion of facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage (Article 5(1)(e)). These prohibitions reflect a broader normative commitment: certain uses of AI are considered inherently incompatible with fundamental rights.

The AI Act’s governance architecture is multi-layered. Each Member State designates national competent authorities and market surveillance authorities (Articles 70–74). At the European level, the AI Office within the European Commission coordinates implementation; the European Artificial Intelligence Board (composed of Member State representatives)

³¹ See the AI Act, Article 6(3); and European Commission, implementing measures on the classification of high-risk AI systems (forthcoming).

provides advice and recommendations; a Scientific Panel of independent experts offers technical guidance; and an Advisory Forum ensures stakeholder engagement (Articles 64–69).

The temporal framework of application is graduated. Chapters I and II (including the prohibitions of Article 5) apply from 2 February 2025. The provisions on general-purpose AI models and key governance-related chapters apply from 2 August 2025. The Regulation is applicable from 2 August 2026, which includes—inter alia—the requirements for high-risk systems listed in Annex III (including judicial AI). Article 6(1) (high-risk AI systems embedded in regulated products) and the corresponding obligations apply from 2 August 2027.³²

For colleagues from Asian jurisdictions present at this workshop, the AI Act represents the European Union’s ambition to set global regulatory standards for artificial intelligence, much as the GDPR has become a de facto international standard for data protection. Whether it will achieve comparable extraterritorial influence remains to be seen. What is certain is that any jurisdiction developing AI tools for its judiciary will find itself engaging—directly or indirectly—with the regulatory categories and normative principles that the AI Act establishes.

10. Italy’s National AI Legislation: Law No. 132/2025

Against that European baseline, Italy has adopted a dedicated national statute that complements the AI Act with provisions tailored to the domestic constitutional and institutional setting.³³

Italy is widely described as the first EU Member State to enact a comprehensive national law on artificial intelligence. Law No. 132 of 23 September 2025, in force since 10 October 2025, complements the EU AI Act with provisions specifically adapted to the Italian legal order.

For the judiciary, the provision of paramount importance is Article 15, which lays down four core rules.

First, it enshrines an absolute *riserva di umanità* for every decision involving the interpretation and application of the law, the evaluation of facts and evidence, and the adoption of judicial measures: “In cases where artificial intelligence systems are used in judicial activity, every decision concerning the interpretation and application of the law, the assessment

³² AI Act, Article 113.

³³ Law No. 132 of 23 September 2025, published in the *Gazzetta Ufficiale* No. 231 of 3 October 2025, entered into force on 10 October 2025.

of facts and evidence, and the adoption of judicial measures is always reserved to the magistrate.”³⁴ The adverb “always” and the adjective “every” are legally significant: the formula is drafted to admit no exception.

Second, Article 15(2) empowers the Ministry of Justice to regulate the use of AI for the organisation of justice services, the simplification of judicial work, and ancillary administrative activities. The structure is deliberately binary: the decisional core (interpretation, evaluation, decision) remains reserved to the human judge; organisational and administrative support functions are, in principle, open to AI deployment under ministerial regulation.

Third, Article 15(3) establishes a transitional authorisation regime: until the full implementation of Regulation (EU) 2024/1689, experimentation with, and deployment of, AI systems in ordinary judicial offices requires prior authorisation by the Ministry of Justice, after hearing the national AI authorities designated by the same law. The provision is notable both for its prudential character and for the institutional tension it may generate—an issue to which this paper will return in the discussion of judicial self-government and executive steering.

Fourth, Article 15(4) links AI literacy to judicial training. In preparing the policy guidelines for magistrates’ training within the institutional framework of the Higher School of the Judiciary (SSM), the Minister of Justice is required to promote teaching activities on AI and on the use of AI systems in judicial work, including both basic and advanced digital training, the sharing of relevant competences, and awareness of benefits and risks. The provision also extends, for the same purposes, to the training of administrative staff.

Beyond Article 15, the statute contains further provisions relevant to legal practice and litigation. Article 13 regulates the use of AI in the legal profession: it frames AI as a tool for instrumental and support activities, with the primacy of the intellectual work that constitutes the professional service, and it requires that clients receive clear, simple, and exhaustive information about the AI systems used by the professional. Article 17 amends Article 9(2) of the Code of Civil Procedure, assigning to the specialised business court divisions jurisdiction over cases whose object is the functioning of an AI system. Article 26 introduces aggravated criminal consequences for offences committed by means of AI systems in specified circumstances and creates a new offence of unlawful dissemination of

³⁴ Article 15(1), Law No. 132/2025.

images, videos, or voices generated or altered through AI systems and capable of misleading as to their genuineness.

The Italian national law operates within a normative ecosystem of remarkable complexity. It must be read together with the AI Act itself; the CEPEJ European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems (2018) and the CEPEJ Guidelines on the Use of Generative AI for Courts (adopted December 2025); the UNESCO Guidelines for the Use of AI Systems in Courts and Tribunals (formally launched in December 2025);³⁵ the Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225, 2024); and the updated Product Liability Directive (Directive (EU) 2024/2853). This multi-level lattice creates a panorama whose practical harmonisation—and, in particular, the allocation of responsibilities among regulators, court administrations, and judicial self-government—remains to be fully defined.

For colleagues from Japan, South Korea, and Taiwan—jurisdictions whose legal systems routinely combine framework legislation with executive decrees and ministerial ordinances—the Italian structure may appear familiar. The interaction between primary legislation, ministerial implementing measures, and advisory or deliberative positions taken by the judiciary’s self-governing body (see the next section) mirrors, in its essential logic, the regulatory architecture of many East Asian legal systems.

11. The Recommendations of the High Council of the Judiciary

Within this multi-layered normative framework, the first internal institutional response came from Italy’s constitutional organ of judicial self-government.

The Italian High Council of the Judiciary (Consiglio Superiore della Magistratura, CSM) is the constitutional body entrusted with the self-governance of the judiciary. Under Articles 104–107 of the Constitution (in particular Article 105), it is responsible for recruitment, assignments, transfers, promotions, and disciplinary measures concerning judges. It is chaired by the President of the Republic and is composed of members elected by the judiciary and by Parliament (with the First President and the Prosecutor General of the Court of Cassation sitting **ex officio**). For

³⁵ UNESCO, *Guidelines for the Use of AI Systems in Courts and Tribunals*, 2025, articulated in 15 principles including transparency, accountability, human oversight, and human rights protection. A striking finding of the UNESCO 2024 survey: only 9% of judicial operators (judges, prosecutors, lawyers) had received AI-specific training, despite 44% already using AI tools and 73% favouring mandatory regulation.

an international audience, the CSM's practical significance lies in its role as guarantor of judicial independence from executive interference.³⁶

On 8 October 2025—two days before Law No. 132/2025 entered into force—the CSM adopted in plenary session its first Recommendations on the use of artificial intelligence in the administration of justice.³⁷ The document sets out, in an expressly non-exhaustive list, twelve categories of AI-assisted activities regarded as *prima facie* permissible in the transitional phase, including: case-law research; legislative and doctrinal research; the drafting of summaries and supporting materials; the organisation of judicial work and statistical reporting; the pseudonymisation of decisions intended for publication; and linguistic revision.

The core normative message is restrictive as to “judicial activity in the strict sense”. Until EU-compliant systems are available, the Recommendations state that unauthorised use of AI systems in adjudication must be excluded; by contrast, they contemplate controlled experimentation in a protected environment—subject to anonymisation, traceability, and the joint supervision of the Ministry of Justice and the CSM—and they stress that any AI output must never become the exclusive or predominant basis of the judge's conviction.

A significant institutional tension emerges when these Recommendations are read alongside Article 15(3) of Law No. 132/2025, which assigns the authorisation of experimentation and deployment of AI systems in ordinary judicial offices to the Ministry of Justice, after hearing the national authorities designated under Article 20. The constitutional question—who governs the digital tools that judges use? —turns, in practice, on the boundary between the Ministry's responsibilities for the organisation of judicial services (Article 110 of the Constitution) and the CSM's mandate to safeguard judicial autonomy and independence (Articles 101 and 104).

The Recommendations also advance several governance proposals: the establishment of institutional technical tables with the Ministry; the creation of a permanent multidisciplinary technical group (possibly including external experts) to conduct independent assessments and develop shared methodological criteria; and, more broadly, a strategic plan—agreed with the CSM—privileging solutions that maximise auditability and data protection. In this light, the document explicitly notes the emergence of locally deployable (often open-source) AI systems that can run large

³⁶ On the CSM, see Articles 104–107 of the Italian Constitution. For an English-language overview, see the European Network of Councils for the Judiciary (ENCJ) database.

³⁷ CSM, Recommendations on the Use of Artificial Intelligence in the Administration of Justice, plenary deliberation of 8 October 2025.

language models on-device without reliance on the cloud, and it treats their controlled testing as a matter of institutional preparedness.

Colleagues from South Korea—where the National Court Administration has launched a Judicial Artificial Intelligence Committee (April 2025)³⁸ to steer AI adoption within the judiciary—and from Japan—where the Supreme Court’s General Secretariat plays a central coordinating role in judicial administration—will recognise in the CSM–Ministry dialectic an instance of a broader governance dilemma: centralised administrative efficiency versus decentralised judicial autonomy.

12. Judicial Training in the Age of AI

A complementary institutional response to the governance tensions outlined above is training. While the CSM claims a role in steering judicial AI, Article 15(4) of Law No. 132 of 23 September 2025 explicitly links the subject to the national planning of judicial education: in preparing the programme guidelines for magistrates’ training under Legislative Decree No. 26/2006, the Minister of Justice must promote teaching activities on AI and its uses in judicial work, aimed at basic and advanced digital training, the acquisition and sharing of digital skills, and awareness of both benefits and risks.

Within that framework, the Higher School of the Judiciary (SSM) has already offered targeted initiatives. A representative example is the webinar course P26026, “Drafting Judicial Decisions in the Algorithmic Society”, scheduled in four sessions (8, 15, 22, and 29 May 2026).³⁹ The course is designed as a practical contribution on sound techniques for “constructing” judicial decisions—capable, after due process, of sustaining the final outcome on appeal—and it explicitly situates those techniques within what it calls the modern algorithmic society.

The programme articulates six thematic axes: (a) it begins with judicial reasoning, focusing on psychological heuristics, the judicial syllogism, the hermeneutic circle, and recurrent “traps” and distortions in both fact-finding and legal qualification, including typical deliberative dynamics; (b) it then turns to judicial interpretation, recalling rules and limits—especially when general clauses and indeterminate legal concepts are at stake; (c) it addresses legal language and the appropriate style of judicial writing; (d) it examines the duty to give reasons and core techniques of motivation, including recurring issues such as *obiter dicta*, and forms of

³⁸ Supreme Court of Korea, press release, 28 April 2025.

³⁹ SSM, Training Programme 2026, Course P26026.

reasoning *ad abundantiam* or *per relationem*, with an eye to identifying both the requirements and the defects of adequate reasoning under appellate scrutiny; (e) it highlights the growing pressure—in the age of AI—towards the formal normalisation of decisions so as to feed databases coherently; the course uses the familiar maxim “garbage in, garbage out” (GIGO) to stress the dependence of any automated processing on the quality and completeness of the input; (f) finally, it treats the prerequisites for building well-structured “merit-level” databases (trial and appellate courts), premised on the idea that well-organised, accessible, and non-duplicative repositories can discourage unfounded appeals and enable the study of interpretive trends across comparable issues.

Points (e) and (f) deserve particular emphasis for present purposes. The prospect that standardisation may become a prerequisite for effective AI deployment raises a structural question about the relation between form and substance in adjudication. GIGO captures only part of the difficulty. The other part is normative and institutional: once machine readability becomes a performance condition, there is an incentive to write in ways that optimise the tool’s output. Over time, that incentive may reshape not only drafting practices but also the cognitive and argumentative habits through which judges stabilise reasons.

The programme also reflects a broader infrastructural asymmetry. Compared to the Supreme Court’s CED/Italgiure environment, merit-level jurisprudence remains far less uniformly curated and retrievable at scale. If AI is to assist responsibly in analysing jurisprudential trends beyond the Supreme Court, the construction of comprehensive, accurate, and genuinely accessible merit-level databases is not an optional add-on but a precondition.

13. From Theory to Practice: National Rollout of Microsoft Copilot 365

On 1 January 2026, the Ministry of Justice enabled Microsoft Copilot 365—integrated within the Microsoft 365 (Office 365) environment already used in judicial offices—as a tool to support the daily work of career judges. The initiative was formally presented, in early January 2026, as an experimental and voluntary programme (with an opt-out mechanism) explicitly confined to support functions and without any intended incidence on the decisional core of adjudication.⁴⁰

⁴⁰ Ministerial note by Antonella Ciriello, Head of the Department for Technological Innovation, 10 January 2026.

The Ministry's framing is significant for the present paper. Copilot is not introduced as a substitute for judgment, but as an instrument of AI literacy under real working conditions: the underlying premise is that, since generative AI is already entering the courtroom ecosystem through lawyers and parties, judges must understand—practically and not merely in the abstract—its limits, failure-modes, and proper use.

Within that approach, the Ministry has emphasised a set of safeguards intended to keep experimentation within a controlled perimeter. The tool is described as operating within the justice domain (i.e., the ministerial Microsoft tenant), with segregation of judicial data and security measures designed to prevent uncontrolled circulation of inputs and outputs. In the same spirit, legal retrieval is expected to be anchored in certified institutional repositories—above all the Court of Cassation's Italgire/ItalgireWeb—so that AI-assisted drafting and research are grounded in quality-controlled sources rather than the open internet.⁴¹

The inaugural webinar of 26 January 2026 is instructive because it clarifies how the administration itself conceptualises the risk profile of generative AI in judicial work. Antonella Ciriello explicitly characterised generative AI as a disruptive tool whose uptake cannot be delayed, yet whose use requires methodical training. She singled out automation bias (the uncritical acceptance of plausible outputs) as a central danger and insisted that AI must remain a support to the magistrate's work rather than a source of heterodetermination of the judicial decision. In her formulation, the premise is both ontological and juridical: *iurisdictio*—the act of “saying the law”—remains the exclusive province of the human intellect.

From the perspective of judicial self-governance, the intervention of CSM Councillor Forziati converged on the same point while sharpening its constitutional edge. He endorsed experimentation only under conditions of traceability, clarity of technical specifications, and effective data-protection guarantees, stressing that the judiciary cannot treat AI as a private user would: the institutional setting requires demonstrable safeguards and a vigilant, critically informed stance. Forziati also articulated the key binary that will reappear throughout this paper: an absolute limit for the decisional nucleus (interpretation and application of law, evaluation of facts and evidence, adoption of judicial measures), coupled with a conditional openness to organisational, documentary, and ancillary uses.

Early feedback from the restricted pilot phase conducted during 2025—reportedly involving selected offices (including, inter alia, the Court of

⁴¹ Details on the technical perimeter and the documentary sources designated as “certified” are drawn from the Ministry's written notice and from the launch webinar of 26 January 2026, both in the author's files.

Cassation, the Court of Appeal of Bari, and the Tribunal of Catania)—points in a predictable direction. The current predictive value of generative AI for judicial outcomes remains limited; its concrete utility lies rather in accelerating case-law research, assisting in the drafting and revision of texts, and supporting the analysis and organisation of documentary materials, always subject to rigorous human oversight and to the requirement of verification against authoritative sources.⁴²

Against this background, the Copilot initiative should be described in measured terms: it is a large-scale, nationwide experimentation with a general-purpose generative AI tool in judicial offices, conducted under an expressly support-only rationale and under a declared commitment to a closed and controlled environment. Those safeguards mitigate—but do not dissolve—three governance questions that remain structurally relevant.

First, governance and constitutional allocation of powers: the Ministry’s choice of a centralised standard tool intersects with the CSM’s responsibility to safeguard judicial independence and with its declared preference for solutions that maximise auditability and institutional control. Even if the tool is confined to support functions, the locus of governance over the digital work-environment of judges is not a neutral administrative detail.

Second, digital sovereignty and vendor dependence: reliance on a single commercial proprietary platform—operated by a foreign corporation—may be consistent with data-protection safeguards, yet it can still generate forms of technological dependency (lock-in), especially when the platform becomes a routine medium of drafting, research, and internal communication.

Third, durability and enforceability of safeguards: closed environment assurances, segregation commitments, and security controls are meaningful only if they remain stable over time, are technically auditable, and can be effectively enforced in practice. In a field where tools, contracts, and model capabilities evolve rapidly, governance must be conceived as an ongoing regime of verification rather than as a one-off ex ante clearance.

The Swiss comparison is illuminating. As Judge Christoph Hurni will detail in his presentation, the Swiss Federal Supreme Court has pursued a radically different path: the development of a sovereign, open-source *Swiss Justice Base Model* (SJB_M), trained exclusively on Swiss legal data using Meta’s Llama architecture, and deployed on-premise through a custom

⁴² The identification of the pilot sites and the internal assessment are based on administrative materials and presentations in the author’s files.

interface (ChatTF) alongside a LibreOffice integration.⁴³ The Swiss model embodies a philosophy of digital sovereignty that stands in contrast to Italy's commercial-platform approach. For the jurisdictions present at this workshop — each confronting comparable choices in its relationship with major technology corporations — the Italy-Switzerland divergence offers a particularly instructive comparative study.

Taiwan offers a further instructive model. Its approach is built on two distinct but complementary tracks. The first is domain-specific and judiciary-controlled: on 27 August 2023, the Judicial Yuan announced an AI-assisted judgment drafting system built on Chunghwa Telecom's TMT5 model, trained on approximately twenty-five years of court judgments (1996–2021), which provides judges — and, prospectively, citizen judges participating in the new lay-judge system — with sentencing recommendations and analogous-case retrieval, while preserving full human decisional authority.⁴⁴ The second track is national and general-purpose: the government-funded TAIDE initiative (*Trustworthy AI Dialogue Engine*), launched by the National Science and Technology Council in April 2023, has developed a sovereign large language model for Traditional Chinese, based on Meta's open-source Llama architecture,⁴⁵ deployed in education, healthcare, and public administration — but not trained on judicial data and institutionally separate from the Judicial Yuan's system. The coexistence of these two tracks illustrates a model of sovereign AI infrastructure in which the judiciary retains full control over its own domain-specific tool while benefiting from the broader national technological ecosystem — a model that repays careful study by any jurisdiction seeking to reconcile digital sovereignty with practical feasibility.

14. AI in Action: PRODIGIT versus CED—Technical Challenges

A final point of continuity with the Copilot discussion should be noted. Copilot is a horizontal, general-purpose tool introduced within the ordinary judiciary; by contrast, Italy's most mature data-driven initiative to

⁴³ Ch. Hurni, "Digital Justice in Switzerland and the Use of AI at the Swiss Federal Supreme Court," *supra*; see also the five ethical principles adopted by the Swiss Federal Supreme Court.

⁴⁴ Judicial Yuan, Press Release, 27 August 2023. On the TMT5 model and the citizen-judge dimension, see the contribution of Professor Kuan-Ling Shen (Jan) to the Kyoto international workshop.

⁴⁵ On TAIDE, see Executive Yuan (R.O.C.), Trustworthy AI Dialogue Engine (TAIDE) (2024), available at <https://english.ey.gov.tw>; the commercial release of TAIDE-LX-7B (based on Llama 2) was on 15 April 2024; an upgraded model based on Llama 3 (TAIDE-LX-8B) followed shortly after.

date has developed within a specialised jurisdiction—tax justice—where the institutional owner of the information system is the Ministry of Economy and Finance rather than the Ministry of Justice.

The PRO.DI.GI.T. project is a programme promoted by the Governing Council of Tax Justice (CPGT) and implemented within the Ministry of Economy and Finance’s tax-justice administration⁴⁶. Its architecture is best understood through its core lines of intervention: (a) the digitalisation of internal CPGT processes and the reengineering of its institutional web infrastructure; (b) the creation of a national, publicly accessible database of tax decisions at the merit level (first and second instance), published in pseudonymised full text; (c) the experimental development of legal-analytics tools, including semantic retrieval and assisted drafting; (d) a pilot track for “predictive” modelling in tax litigation; and (e) a distributed digital laboratory for the tax judge (TribHub), activated on an experimental basis in multiple regions.⁴⁷

A pivotal concrete outcome is the public database of tax-merit decisions, launched in 2024 and progressively updated, which already makes available a very large corpus of native-digital judgments (with full-text access, search filters, and downloadable PDFs). The institutional significance of the database is twofold. First, it reduces the asymmetry previously produced by the fact that the tax administration had privileged access to comprehensive merit-case material. Second, it provides the indispensable substrate for any serious form of legal analytics in tax litigation—while simultaneously foregrounding the governance problem of representativeness and completeness, since a dataset that is not exhaustive (or not evenly representative across courts) may bias both descriptive statistics and any downstream AI-assisted analysis.

In parallel with the database, PRODIGIT has also served as a testbed for AI-supported summarisation and structured information extraction. The technical literature produced around the project describes experiments combining extractive and abstractive approaches, including LLM-based summarisation evaluated by domain experts, with explicit attention to the risk-profile of hallucinations and omissions in complex reasoning. In this setting, summarisation is not a merely rhetorical aid: it is conceived as a

⁴⁶ The ICT component is entrusted to SOGEI (the Ministry’s in-house IT company). It is financed under the PON “Governance and Institutional Capacity 2014–2020” through REACT-EU resources (approximately €8.34 million). The PON is a multi-year investment plan funded by the European Union and the Italian government to reduce economic and social disparities. It focuses on modernizing public administration, enhancing efficiency, and increasing transparency.

⁴⁷ I. Ferranti, “Prodigit, come funziona il progetto per la giustizia tributaria digitale”, *Agenda Digitale*, 2023: project structure, TribHub, funding figure, links to PNRR context.

way of rendering the factual background, the legal issue, and the decision criteria legible at scale, and therefore searchable and comparable across a national corpus.⁴⁸

By comparison, AI work at the Supreme Court of Cassation must confront a different set of constraints. The Court's documentary infrastructure (CED/Italggiure) gives it an advantage in curated sources, controlled indexing, and institutional responsibility; yet the linguistic and argumentative density of Supreme Court reasoning—its long hypotactic structure, technical terminology, and pervasive internal cross-references—poses well-known challenges for current NLP systems, especially where the task is not mere retrieval but the reconstruction of an interpretive framework (mapping the jurisprudential and doctrinal constellation around a contested point of law). This is also why document pre-processing becomes a governance-relevant technical layer: the more reliable the segmentation, anonymisation, and structural parsing of judgments, the less pressure is placed on downstream generative components.

Recent research based on Supreme Court judgments has therefore explored document-layout analysis as a preprocessing step, using computer-vision models to segment and classify the structural parts of decisions before NLP is applied. These pipelines do not, by themselves, solve the epistemic problem of legal reasoning; but they can materially improve the quality and auditability of the textual inputs on which any legal-analytic or generative tool depends.⁴⁹

15. Pilot Projects Across the Italian Judiciary

⁴⁸ T. Dal Pont et al., "Legal Summarisation through LLMs: The PRODIGIT Project" (arXiv:2308.04416), develops and evaluates the use of large language models for the automated summarisation of Italian tax-law decisions, thereby illustrating the practical applicability of AI within Italian digital justice. A. Fidelangeli et al., "The Summarization of Italian Tax-Law Decisions: The Case of the PRODIGIT Project" (2025) introduces a hybrid AI approach to automatically summarize complex legal rulings. Utilizing large language models and a structured, combined approach, the method improves upon traditional summarization techniques for legal professionals (read the full paper at IEEE Xplore). A. Voza, "Giurisprudenza tributaria di merito: on line la banca dati gestita dal MEF", *Il fisco* 35/2024

⁴⁹ M. Marulli, G. Panattoni, M. Bertini, "A document processing pipeline for the construction of a dataset for topic modelling based on the judgments of the Italian Supreme Court", arXiv:2505.08439 (layout analysis using YOLOv8x as part of a preprocessing pipeline). ArXiv is a free distribution service and an open-access archive for nearly 2.4 million scholarly articles in the fields of physics, mathematics, computer science, quantitative biology, quantitative finance, statistics, electrical engineering and systems science, and economics.

Beyond the nationwide Copilot deployment in the ordinary judiciary and the specialised PRODIGIT initiative in tax justice, several localised pilot projects illustrate the diversity of institutional approaches within Italy.

At the Tribunal of Catania, an experimental collaboration with the National Interuniversity Consortium for Informatics (CINI), under the direction of Mariano Sciacca, President of Division at the Tribunal of Catania, launched the Lexintel pilot, centred on LLM/NLP applications to judicial documents and, in particular, on the design and testing of structured prompts. The experimentation is conducted retrospectively on already concluded civil cases, after anonymisation, and is expressly framed as a quality-and-risk assessment under strict human oversight. It combines document search and synthesis, information extraction and classification, comparison of texts, the construction of internal links across the materials of a file, and the evaluation of complex prompts.⁵⁰ Catania was identified by the Ministerial Observatory as one of the two pilot centres — alongside Bari — for evaluating prompt effectiveness and hallucination-reduction strategies.

In remarks delivered at the ANM webinar on artificial intelligence (12 February 2026)⁵¹, Mariano Sciacca urged a “semantic hygiene” that demystifies AI: generative systems should not be anthropomorphised as “intelligent” but understood as statistically driven agents lacking semantic comprehension. He stressed that, in justice, the central risk is not the “robot judge” but a grey zone of auxiliary uses that can quietly capture the decisional process: opacity (“black box”) undermines reasons-giving; biased or “dirty” training data reproduce and amplify distortions; and hallucinations—systems’ propensity to supply plausible but false answers—make rigorous human verification indispensable. On Sciacca’s account, cognitive offloading and automation bias may de-skill judges over time, standardise thought, and create interpretive “echo chambers”, especially where efficiency incentives privilege machine-readable forms. He therefore advocated a governance stance of constitutional precaution: a robust *riserva di umanità*, intensive training (including disciplined prompting practices), and—crucially—judicial digital sovereignty, favouring tailored, internally governed models trained on controlled legal corpora to avoid structural lock-in to external commercial platforms.

At the Court of Appeal of Bari, judge Nicola Morgese developed an AI agent for querying judicial archives, including a tool that extracts all statutory references from a judicial document and links them via hypertext

⁵⁰ The Lexintel project was publicly launched on 25 October 2025 as a collaborative AI initiative also involving the Bar Association of Catania, with support from the Italian Ministry of Justice, s. “Il Dubbio”, 25 October 2025.

⁵¹ Summary transcript on file with the author.

to the official legislative database (Normattiva). The evocative historical reference to Justinian’s ambition of placing the laws *in promptu* — “at hand” — is suggestive of the continuity between the ancient quest for legal accessibility and its modern digital incarnation.

At the *Tribunal of Siracusa*, Judge Alfredo Spitaleri explored the integration of AI with the electronic civil procedure system (PCT) through structured prompting techniques for the extraction and analysis of procedural documentation.⁵²

In Emilia-Romagna, the regional ER4Justice programme includes the University of Bologna’s AI4Justice laboratory, which has pursued applied research on the implications of AI in courts and has experimented with integrating heterogeneous legal documents and semantic-web elements to improve retrieval, contextualisation, and explainability. In parallel, the AI4Justice research agenda has emphasised the role of interoperable markup standards for legal documents—such as Akoma Ntoso/LegalDocML and LegalRuleML—as a prerequisite for reliable downstream analytics.⁵³

Taken together, these initiatives converge—at least at the level of declared method—on a common premise: AI should operate as a support tool anchored in closed, verified, and institutionally controlled corpora (anonymised case files, curated archives, annotated legal datasets), rather than as an open-ended system drawing from the unfiltered internet. The aim is not to deny the usefulness of generative models, but to discipline their use by constraining inputs and audit trails, thereby mitigating hallucinations and strengthening accountability.

16. Emerging Case Law: Italian Courts Confront AI

The local pilots discussed above have been accompanied by an equally swift judicial response: Italian courts are already articulating, case by case, the minimum procedural and institutional conditions under which AI-assisted drafting and algorithmic decision-making can remain compatible with the core requirements of adjudication.

⁵² On Bari and Siracusa pilots: “Giustizia 4.0 – Diritti, Algoritmi e Nuovi Equilibri”, Capri conference programme (4 Oct 2025) (listing: Nicola Morgese, Court of Appeal of Bari; Alfredo Spitaleri, Tribunal of Siracusa).

⁵³ On AI4Justice deliverables and markup standards: University of Bologna/ALMA-AI, AI4Justice press release (2021).

A first body of case law is thus emerging on AI-related risks in adjudication and litigation. Five decisions (and lines of authority) merit particular attention.

First, the Court of Cassation parenthetically cautioned that while digital tools may facilitate the drafting of judicial decisions, they also heighten the risk—now “exponentially” increased by the advent of AI—that a judge may draw “aliunde” the reasons for deciding, thereby abdicating the duty to provide an “ineliminable and irreplaceable” evaluative contribution and undermining the very core of judicial thirdness and impartiality.⁵⁴

Second, in a different procedural posture, the Court of Cassation quashed an appellate judgment upholding a conviction whose reasoning relied on “principles of law not affirmed” by the Supreme Court and on precedents “incorrect in the number reported”. It added that the defence had been unable to locate the decisions cited by the appellate court despite searches “also at the offices of the Supreme Court”. Although the judgment does not turn on AI as such, the fact-pattern closely tracks a central risk of AI-assisted drafting in adjudication and advocacy: the insertion, into the reasons for decision, of apparently authoritative—but in fact non-verifiable—citational material.⁵⁵

Third, the Tribunal of Turin (Labour Division, 16 September 2025) applied Article 96(3)–(4) of the Code of Civil Procedure, ordering the claimant to pay €500 to each defendant and €500 to the fine fund, after finding that the action had been brought in bad faith or (at least) with gross negligence. The statement of claim—explicitly described as drafted “with the support of artificial intelligence”—was characterised as an incoherent accumulation of abstract, largely irrelevant statutory and case-law citations, unsupported by concrete factual allegations.

Fourth, the Tribunal of Florence (14 March 2025) addressed the practical problem of fabricated or non-existent citations (“hallucinations”) in a counsel’s procedural submissions allegedly drafted with ChatGPT and considered the consequences that such filings may have (or may fail to have) under the ordinary procedural toolbox, including the potential relevance of Article 96 of the Code of Civil Procedure.

Fifth, in administrative law, the Council of State has developed—through a set of decisions arising from algorithmic management of teacher assignments under Law no. 107/2015—a comparatively mature framework often summarised as a domestic standard of “algorithmic legality”. Its

⁵⁴ Court of Cassation, Criminal Division, Section III, Ruling No. 34481/2025, 22 October 2025.

⁵⁵ Court of Cassation Criminal Division, Section III, ruling no. 25455/2025, 10 July 2025.

recurring pillars are: (a) knowability (the right to know the algorithm and its relevant features); (b) non-exclusivity of automated decision-making (the irreducible requirement of human oversight and responsibility); and (c) non-discrimination (the obligation to prevent, detect, and mitigate discriminatory effects).⁵⁶ For colleagues from jurisdictions where algorithm-driven public-administration decisions have generated acute controversy—e.g. the Netherlands’ SyRI litigation on algorithmic welfare-fraud detection—this Italian “algorithmic legality” framework may offer a productive comparative reference point.

Taken together, this emerging case law sends a clear signal: Italian courts are already enforcing, through ordinary constitutional and procedural reasoning, robust requirements of human responsibility, verifiability, and judicial impartiality that anticipate the governance logic of the EU AI Act—well before the Regulation becomes generally applicable (2 August 2026) and even before the full application of the core high-risk regime (2 August 2027).

17. Comparative Remarks: Four Structural Tensions

The Italian experience with artificial intelligence in the civil justice system may be distilled into four structural tensions that crystallise the challenges facing any jurisdiction at the intersection of judicial independence, technological innovation, regulatory governance, and epistemological integrity.

(a) Executive Influence versus Judicial Independence

Article 15(3) of Law No. 132/2025 vests the authorisation power for AI experimentation in the Ministry of Justice — an organ of the executive. The CSM claims a “co-primary role” grounded in the constitutional guarantee of judicial independence (Articles 101, 104, and 108 of the Constitution). The underlying question — *who governs the digital tools that judges use?* — is not merely administrative. If the executive controls the AI systems that structure judicial research, suggest analytical frameworks, or draft textual elements, the independence of the judiciary may be subtly but effectively compromised, regardless of formal safeguards. This tension resonates with debates in other jurisdictions: in South Korea, where the Judicial Artificial Intelligence Committee was established in April 2025 to coordinate AI governance from within the judiciary;⁵⁷ in

⁵⁶ Council of State, Rulings Nos. 2270/2019 and 8472/2019.

⁵⁷ Supreme Court of Korea, establishment of the Judicial Artificial Intelligence Committee, press release, 28 April 2025.

China, where the Supreme People’s Court’s centralised Smart Courts programme pursues both efficiency and standardisation objectives under centralised judicial governance;⁵⁸ in the United States, where the federal courts’ relationship to commercial legal-technology providers raises analogous questions; and in Latin America, where — as discussed above — the experiences of Brazil, Colombia, and Mexico illustrate that the human-in-the-loop safeguard is not a luxury of consolidated democracies but a universal institutional necessity.

(b) Proprietary Platform Centralisation versus Digital Sovereignty

The decision to deploy Microsoft Copilot 365 as the nationwide judicial AI tool — rather than developing a sovereign, open-source solution — creates a vendor dependency without precedent in any European judicial system. The CSM’s expressed preference for on-premise and open-source solutions reflects a concern for digital sovereignty that finds its most complete realisation in the Swiss model: the Federal Supreme Court’s *Swiss Justice Base Model* (SJBM), built on Meta’s open-source Llama 3.x architecture and trained exclusively on Swiss legal data, represents a conscious choice to maintain full institutional control over the AI infrastructure that serves the judiciary.⁵⁹ Taiwan’s dual-track model — combining a judiciary-controlled domain-specific tool with a national sovereign language model, as analysed earlier in this paper — offers a further alternative in which institutional autonomy and technological sovereignty are pursued simultaneously. The Italian choice, while defensible on grounds of speed and practicality, carries risks that will require sustained institutional vigilance.

(c) The Multi-Level Regulatory Overload

The normative ecosystem governing AI in the Italian judiciary comprises: the EU AI Act; the Council of Europe Framework Convention (CETS No. 225); the CEPEJ Ethical Charter and Guidelines on Generative AI; the UNESCO Guidelines; Law No. 132/2025; the CSM Recommendations; forthcoming ministerial implementing decrees; and the Data Protection Authority’s interpretive guidance. This multi-level lattice generates a risk of regulatory uncertainty and *compliance fatigue* — a condition in which the sheer complexity of the normative framework paralyses rather than guides implementation. As the United States experience suggests —

⁵⁸ Supreme People’s Court of the People’s Republic of China, *Opinions on Regulating and Strengthening the Applications of AI in Judicial Fields*, December 2022; see also D. Xiuming, speech at the 20th Meeting of Chief Justices of Supreme Courts of SCO Member States, Hangzhou, 23 April 2025.

⁵⁹ Ch. Hurni, “Digital Justice in Switzerland and the Use of AI at the Swiss Federal Supreme Court,” *supra*.

where, as Professor Lubbers will discuss, the Supreme Court remains “still operating in the analog age” partly because of the absence of a clear regulatory framework — the extremes of over-regulation and under-regulation both carry costs.⁶⁰

(d) Epistemological Transparency versus Adjudicative Autonomy

A fourth tension, less immediately visible but intellectually fundamental, runs through the entire analysis. As argued earlier in this paper, the distinction between questions of fact and questions of law — the boundary that legitimates the very architecture of appellate review, and particularly the jurisdictional identity of a supreme court of cassation — cannot be drawn with mathematical precision. It reflects, ultimately, a choice of judicial policy: how decisional authority should be allocated between the court below and the court of review.

AI intervenes precisely at this fault line. When natural-language processing tools analyse the reasoning of lower-court decisions, identify “merely apparent” motivations, or classify appeals as raising genuine issues of legal principle rather than disguised requests for factual re-examination, they do not apply a pre-existing boundary between fact and law — they *constitute* it operationally. The algorithm’s classification becomes, in practice, a powerful suggestion as to where the line should fall in any given case. This is not a neutral act of retrieval; it is an epistemological intervention that quietly reshapes the scope of cassation review itself.

The philosophical framework developed in the first part of this paper illuminates why this tension is so difficult to resolve. If, as Gehlen argued, the human subject is a *Mängelwesen* whose cognitive architecture is constitutively shaped by its tools — if, as Sini has demonstrated, each great writing practice does not merely transmit thought but generates the thinking subject — then the judge who works with AI-assisted case-law analysis is not performing the same cognitive operation as the judge who worked without it. The horizon of relevance shifts; patterns of reasoning reorganise themselves around the affordances of the tool. The risk is not that AI will “decide” questions of fact (a prohibition that every jurisdiction represented at this workshop has clearly established), but that it will silently redraw the epistemic map on which the fact/law distinction depends, thereby altering the distribution of adjudicative power between levels of jurisdiction without any deliberate normative choice to that effect.

⁶⁰ J. Lubbers, “The U.S. Supreme Court — Still Operating in the Analog Age,” presentation delivered at the Kyoto international workshop.

For the Italian Court of Cassation — whose institutional identity since the nineteenth century has been built upon the claim to review law, not fact — this tension is existential. But it is equally relevant for any supreme court whose competence is defined by the fact/law boundary: from the Japanese Supreme Court's distinction between appeals on points of law and those on procedural grounds,⁶¹ to the Korean Supreme Court's analogous jurisdictional architecture, to the Swiss Federal Supreme Court's delimitation between *Rechtsfragen* and *Tatfragen*. The deployment of AI tools that operate across this boundary — classifying, filtering, suggesting — demands a new form of epistemological transparency: a clear and auditable account of how the tool's design choices interact with the normative allocation of decisional authority.

These four tensions — independence versus control, sovereignty versus convenience, regulatory complexity versus regulatory clarity, and epistemological transparency versus adjudicative autonomy — define the landscape in which the Italian judiciary navigates its "first steps" in the age of artificial intelligence. The Italian experience constitutes, in a sense, a real-time laboratory whose lessons — both achievements and unresolved difficulties — may illuminate the global debate and, in turn, benefit from the comparative insights that a workshop such as this one is designed to generate.

18. Closing Remarks

I wish to close on a note of optimism. My fundamental attitude toward AI in judicial work is constructive and positive. Technology is, as it has always been, an extraordinarily powerful extension of human cognitive capacities that, if wisely governed, can enhance the quality and efficiency of justice—provided that the act of judging, in all its complexity, fallibility, and irreducible humanity, remains the exclusive province and responsibility of the human person, albeit historically and materially mediated by the technological tools and institutional practices of each era.

I speak not only as a former professor but as a sitting justice. I can already imagine the quiet transformation that AI will bring to my daily work at the Court of Cassation. I am aware that, if I research case law through the AI-enhanced CED database, the patterns surfaced by algorithmic retrieval will subtly shape the horizon of what I consider relevant—and that the precedents it does not surface may silently disappear from judicial memory; if I examine a case file that has been triaged through an

⁶¹ See Code of Civil Procedure of Japan, Articles 311 ff. (*jōkoku* appeal). On the Japanese Supreme Court's approach to AI, see the contribution of Judge Uga to this workshop.

accelerated procedure under Article 380-bis of the Code of Civil Procedure, I know that an upstream classification — whether performed by a human or, prospectively, by a machine — has already framed the boundaries of the judicial question before I have read a single page. What is gained is substantial: speed, consistency, access to a breadth of jurisprudential material that no individual mind could survey. What risks being lost is equally real: the serendipity of unexpected discovery, the unstructured reflection that sometimes leads a judge to see a case differently from all its predecessors, and — most fundamentally — the acute awareness that every act of judgment is an act of personal responsibility that no tool can share or absorb.

These are not theoretical concerns. They are the daily texture of adjudication in the age of artificial intelligence. And they converge with the deepest tension identified in this paper: the judge whose horizon of relevance is shaped by algorithmic retrieval is not performing the same cognitive operation as the judge who worked without it. If the philosophical analysis developed in the first part of this paper is correct — if the human subject is constitutively shaped by its tools — then the *riserva di umanità* cannot be understood as the preservation of a fixed, pre-technological human essence.

There is no such essence to preserve.

What *can* be preserved — and what *must* be preserved — is something more precise and more demanding: the structure of personal responsibility. The judge who signs a judgment must be able to account for every step of the reasoning that leads to the operative part — not because the reasoning was produced in pristine isolation from technological mediation (it never was, not even in the age of the printed code and the handwritten opinion), but because the judge has made it *her own* through an act of critical appropriation, verification, and evaluative judgment that no algorithm can perform or simulate. The *riserva di umanità*, properly understood, is not a reservation of biological nature against the machine. It is a reservation of *authorship* — the irreducible demand that a human person stand behind the judgment, not as the operator of a system, but as the author of a reasoned decision for which she bears full personal and institutional accountability.

And it is precisely because I find the balance, on the whole, favourable — because I believe that wisely governed technology enhances rather than diminishes the quality of justice — that I insist on the necessity of what this workshop embodies: sustained international dialogue, rigorous normative analysis, and the institutional courage to embrace innovation while maintaining an unshakeable commitment to the principles of human dignity, judicial independence, and the rule of law.

