

Access to justice before independent and impartial courts in the light of new technologies

Goran Selanec S.J.D.
19 June 2026

Solve the task:

Filipe was born in Belgium. After primary school his parents moved for work to Portugal. They all reside in Portugal now. He cares about his local community a great deal and is very engaged in communal public life.

There are local elections approaching. He is considering engaging as a candidate for local council before the Youth Party of Portugal. However, he was told that he does not have voting rights since he is Belgian.

What are his election rights?

Rule-by-Rule, Step-by-Step

- Is Filipe allowed to vote?
 - **Legal Fact: Election type**
 - Empirical fact: local election in Portugal
 - Empirical fact: Portugal is EU MS
 - **Legal fact – nationality and residence**
 - Empirical fact: Belgian
 - Empirical fact: Belgium is EU MS
 - Empirical fact: residence over three year
 - **Legal fact: Legal Age**
 - Empirical fact: 16 years of age
 - Empirical fact: Belgium set voting age at 16
 - Empirical fact: Portugal set voting age at 18

Opus 4.8 Summary: Filipe's Electoral Rights

Right	Filipe's Position	Conditions
Vote in local elections	Yes	Register before 60-day deadline; residence in Portugal
Stand as candidate (council member)	Yes	Registration + party list inclusion; declaration of non-dual candidacy
Stand as Presidente da Câmara	Likely No	Reserved for Portuguese nationals under permitted derogation
Vote in Belgian local elections simultaneously	No	Must choose one jurisdiction; cannot exercise rights in both
Vote in EP elections in Portugal	Yes	Can choose to vote in Portugal rather than Belgium

VII. Practical Recommendation

Filipe should immediately:

1. Register on the electoral roll at his local Junta de Freguesia before the 60-day deadline (for the 12 October 2025 autárquicas, that was approx. 11 August 2025).
2. Contact the Youth Party of Portugal to confirm his place on the candidate list — the party must formally submit the list to the competent electoral court.
3. Make the required declaration that he is not exercising simultaneous electoral rights in Belgium.
4. Clarify with the party which position on the list he would occupy, keeping in mind the structural limitation on the mayoralty.

The advice he received was not merely incomplete — it ran contrary to binding EU primary law, EU Charter rights, and the implementing Portuguese legislation. As a Belgian EU citizen legally resident in Portugal, Filipe enjoys the full range of local electoral rights, both active and passive, subject only to the procedural registration requirements applying equally to all electors.

Copilot – 2nd Take

- how do you know Filipe is of proper voting age?
 - From the information provided in the document, it is mentioned that Filipe is **16 years of age 1**. However, it's important to note that while Belgium sets the voting age at 16, Portugal sets the voting age at 18 **1**. Therefore, Filipe would need to be at least 18 years old to have voting rights in Portugal.
 - If you have any more questions or need further assistance, feel free to ask! 😊
 - **1**: Slide 3/72 of the presentation.

The Original Advice: Right Conclusion, Wrong Reason

Filipe was told he has no voting rights. That turned out to be correct — but for the entirely wrong legal reason.



REASON GIVEN

**"He cannot vote
because he is Belgian"**

✗ Legally incorrect



CORRECT REASON


**"He cannot vote
because he is 16"**

✓ Portuguese Constitution, Art. 49(1)

⚠ Why it matters: Accepting the wrong reason could permanently exclude Filipe even after he turns 18 — not knowing Belgian EU citizens enjoy full local electoral rights in Portugal upon reaching majority.

Increased Efficiency and Productivity

- AI automates tasks like document review and legal research.
- Streamlined workflows through integration with existing software.
- Faster research by quickly analyzing large legal databases.



Reasoning Machine

- Do DL AIs reason?
- Yes. Not in human manner.
 - Do not to reproduce legal reasoning!
- Discern structures - patterns and distributions - of the input data and reconfigure their computing processes incorporating them into their expanding data base (base of knowledge)
- Identify correlations between the different parameters of case on the input side (lexical groups related to facts and applicable law) and infer one or more models to forecast the likelihood of a specific outcome (lexical groups related to ruling)
 - e.g. in a divorce dispute: identify correlations between (lexical groups related to) the length of marriage, the income of spouses, the existence of adultery on the input side and by adjusting their weights in correspondence to acquired past knowledge it can generate (lexical groups) related to assessed amount of the benefit pronounced on the output side

Dominant models in Legal Practice

- Great majority of AI tools used in legal practice are NLP (natural language processing) AIs, more precisely a sub-set known as LLM (large language models) AIs.
- LLMs are AI systems that use vast training datasets of text to recognize and extrapolate relations and patterns within text based on which it can complete a task of generating appropriate responses to user's instructions (prompts) expressed in human language.
- They rely heavily on syntax and vector mathematics.
- By using different algorithms they can be trained for different task of language manipulation such as summarizing information, identifying relevant sections, pointing out key points, classifying information into various patterns, providing citations and references, finding related documents, generating specific answers to questions. A user may thus give instruction requesting an overview of case-law in a specific area of law classified in different sets such as jurisdictions, courts or even individual judges.

Domain-Specific Embeddings

*Why the map inside an AI determines
whether it finds the right law — or the wrong one.*

For lawyers who need to understand the tools they're using

Mapping

- **Washington DC** is located at 38.9 degrees North and 77 degrees West. We can represent this using a vector notation:
 - **Washington DC** is at [38.9, 77]
 - **New York** is at [40.7, 74]
 - **London** is at [51.5, 0.1]
 - **Paris** is at [48.9, -2.4]
- This is useful for reasoning about spatial relationships. You can tell **New York** is close to **Washington DC** because 38.9 is close to 40.7 and 77 is close to 74. By the same token, **Paris** is close to **London**. But **Paris** is far from **Washington DC**.

JUSTICE

- [0.3058549463748932, 0.36481973528862, -0.2223075032234192, 0.051489319652318954, 0.2177533358335495, 0.6139580607414246, 0.20636603236198425, -0.1489017903804779, -0.023582270368933678, 0.39647945761680603, 0.5525206923484802, 0.12038766592741013, -0.4224829077720642, -0.011456784792244434, -0.10493946820497513, -0.25726836919784546, -0.055846814066171646, -0.21123136579990387, -0.037917956709861755, 0.026921391487121582, -0.27512127161026, -0.14914382994174957, 0.3690960705280304, -0.1340155303478241, 0.21202640235424042, -0.22819402813911438, -0.2585029900074005, 0.4864526391029358, -0.24878238141536713, 0.14480730891227722, -0.4541610777378082, 0.4139264225959778, -0.1055862233042717, 0.5027562379837036, -0.23430921137332916, 0.13076603412628174, -0.1671728640794754, 0.0961972177028656, 0.010481687262654305, -0.35954153537750244, 0.11985704302787781, -0.27856573462486267, 0.40210551023483276, 0.021673021838068962, -0.1790880560874939, 0.01745273917913437, 0.00843943003565073, -0.16605590283870697, -0.22320057451725006, 0.2120678573846817, 0.2966324985027313, 0.1291898488998413, -0.2620503604412079, 0.5413062572479248, -0.58146733045578, -0.31027501821517944, -0.24212610721588135, -0.08700574934482574, 0.2873508036136627, 0.13863641023635864, -0.1177591010928154, -0.3583690822124481, 0.05888424068689346, -0.15392374992370605, 0.46106770634651184, 0.07916717976331711, 0.108493871986866, 0.17081554234027863, 0.3549806475639343, 0.24713541567325592, 0.14140073955059052, -0.4236290752887726, 0.03451409190893173, -0.0776522234082222, -0.11378617584705353, -0.5364951491355896, -0.10805702209472656, 0.32199516892433167, -0.12212540209293365, -0.4539981484413147, -0.050136443227529526, 0.628773033618927, -0.14440518617630005, -0.2920602262020111, 0.22030875086784363, 0.13635806739330292, 0.00844517257064581, -0.17004585266611328, -0.08113452792167664, -0.4420273005962372, 0.028931530192494392, 0.040615614503622055, 0.08454050868749619, -0.12730541825294495, -0.16914606094360352, -0.12474681437015533, -0.4694853127002716, 0.6898261904716492, 0.19227726757526398, 0.16400161385536194, 0.10631895810365677, -0.5279147624969482, -0.10503244400024414, -0.22589358687400818, 0.13231337070465088, 0.22063878178596497, 0.04175635427236557, 0.012677093036472797, -0.06494039297103882, -0.19329488277435303, -0.22649024426937103, 0.09435412287712097, -0.3088231682777405, -0.009853686206042767, 0.12043270468711853, -0.31781548261642456, -0.258686363697052, -0.36225223541259766, 0.18566477298736572, -0.4277484118938446, 0.12490992993116379, -0.263084352016449, -0.09273895621299744, -0.03705182299017906, -0.34295567870140076, 0.1452440619468689, -0.1395331174135208, -0.0924496203660965, -0.3243474066257477, 0.13045518100261688, -0.08529586344957352, 0.20504112541675568, -0.0684538185596466, -0.22366493940353394, 0.3995629847049713, -0.21737302839756012, -0.03397352248430252, 0.04532643035054207, -0.02849491313099861, -0.22167456150054932, 0.4076932966709137, -0.26372864842414856, 0.19332173466682434, -0.4041522443294525, -0.028065267950296402, -0.19069507718086243, 0.1760089248418808, 0.33836162090301514, 0.07879278063774109, 0.06858856976032257, -0.27528393268585205, -0.787016749382019, 0.2446596324443817, 0.18336167931556702, 0.0217819232493639, -0.09623502939939499, 0.08253531157970428, -0.24581566452980042, 0.1798137128353119, 0.3183457553386688, -0.02165018394589424, 0.2959766387939453, 0.3761600852012634, 0.06308282911777496, -0.7761154174804688, -0.4726817011833191, 0.6156711578369141, -0.13821038603782654, 0.240960493683815, -0.12212163209915161, 0.450241357088089, -0.003809584304690361, 0.09408561885356903, 0.10472561419010162, 0.30536749958992004, -0.18514318764209747, -0.25358423590660095, -0.06180192530155182, -0.16169263422489166, 0.4035210609436035, 0.2206708937883377, 0.3564136326313019, -0.2744482159614563, -0.08996666222810745, 0.012530840002000332, 0.20459653437137604, -0.023988448083400726, 0.33522719144821167, 0.09851924329996109, 0.02540556713938713, -0.15383997559547424, -0.20524699985980988, -0.273471474647522, -0.18816334009170532, -0.1168024092912674, 0.1542869210243225, -0.2566792666912079, 0.10040758550167084, 0.2978368401527405, 0.24682076275348663, 0.22066274285316467, 0.031800124794244766, -0.2885070741176605, 0.09910909086465836, -0.12450578063726425, 0.1971411406993866, -0.38965919613838196, 0.3671571612358093, 0.34351563453674316, 0.16509993374347687, -0.041634369641542435, 0.2064838856458664, 0.13155871629714966, 0.44245702028274536, -0.002932378090918064, 0.10958034545183182, -0.32727187871932983, 0.26840487122535706, 0.09633728861808777, 0.8179660439491272, 0.23226137459278107, -0.44188815355300903, 0.147987499833107, -0.16952356696128845, 0.1762319654226303, 0.41069719195365906, -0.3585379421710968, 0.08348345011472702, -0.6487800478935242, 0.03170100599527359, 0.03463808447122574, -0.1129198670387268, -0.06297271698713303, 0.03381196781992912, -0.4863276481628418, 0.45702308416366577, 0.08004499226808548, -0.4242900013923645, 0.4159078598022461, -0.1447920799255371, 0.15365716814994812, -0.1172354444861412, -0.06826122850179672, 0.16746088862419128, 0.19755375385284424, 0.27143803238868713, 0.045931749045848846, -0.6983194947242737, 0.020827794447541237, 0.25797271728515625, 0.429937481880188, 0.16930320858955383, 0.2867017686367035, 0.06390050053596497, 0.031147923320531845, 0.29502665996551514, 0.3084077835083008, 0.04922319948673248, 0.16442032158374786, 0.34734466671943665, 0.1892312616109848, -0.08626750111579895, -0.06930122524499893, 0.5451787114143372, -0.05609678477048874, 0.027657240629196167, -0.5087507963180542, -0.11402104794979095, 0.42347797751426697, -0.13967163860797882, 0.3406089246273041, 0.2660832703113556, -0.18712171912193298, -0.16602614521980286, 0.3034486770629883, -0.099801205098629, 0.08338838815689087, 0.030712855979800224, 0.03674490004777908, -0.07732322812080383, 0.13995741307735443, -0.3316799998283386, -0.05812756344676018, 0.22520290315151215, -0.11088104546070099, 0.07114162296056747, 0.15775351226329803, 0.08936847746372223, 0.17194586992263794, 0.10405148565769196, -0.13586091995239258, -0.22330242395401, -0.3082481026649475, -0.06716524064540863, -0.043051812797784805, -0.0527329221367836, -0.1076473817229271, 0.033714067190885544, -0.5254990458488464, -0.3826981782913208]

What Is an Embedding?

The Core Idea


Computers speak numbers. Legal texts speak words. An embedding is the translation — it converts words, sentences, or whole documents into a list of numbers (a vector) placed in a vast mathematical space.

The rule that governs this space:

Similar meaning → numbers close together

Different meaning → numbers far apart

The Semantic Map



Why It Matters: The Two Librarians

*When you query a legal AI, the embedding model decides which documents the AI reads before answering.
Get the wrong librarian — get the wrong answer.*



General-Purpose Embedding

Trained on: Wikipedia, news, novels, Reddit...

You ask: "Find cases about consideration in contract formation"

It retrieves: Philosophy texts on 'careful deliberation'...
newspaper editorials about 'due consideration'...

Because in general English, that's what 'consideration' means.



Legal Domain Embedding

Trained on: 20B+ tokens of case law, statutes, contracts...

You ask: "Find cases about consideration in contract formation"

It retrieves: Promissory estoppel cases, Restatement sections, offer & acceptance doctrine...

Because that's what 'consideration' means in law.

Why Legal Language Breaks General AI

Legal language is a dialect. Four ways it systematically defeats general-purpose embedding models:

1 Terms of Art

"Consideration" = exchange of mutual promises in law

"Leave" = procedural authorisation in litigation

"Stay" = suspension of proceedings

General AI has learned the ordinary meaning — it maps to the wrong concept.

2 Citation Structure

"Donoghue v Stevenson [1932] AC 562"

This opaque string carries enormous legal weight (foundational negligence duty of care). A general model treats it as an unreadable proper noun. A legal model knows exactly what it means.

3 Procedural Vocabulary

"Filing", "pleading", "discovery", "judgment", "costs" — all have specific procedural meanings that are different from ordinary usage.

General models average across all uses. Legal models learn the procedural sense.

4 Jurisdiction Specificity

"Tribunal" means different things in EU law, UK administrative law, and Singapore procedural law.

General models produce a blurred, jurisdiction-agnostic average. Domain-trained models learn context.

So What?

01

Fluent ≠ Correct

A legal AI can produce a beautifully written, citation-rich answer based on the wrong retrieved documents. The generative model was doing its job — the embedding model sent it to the wrong part of the library. You cannot tell from the output alone.

02

Jurisdiction Gap Is Real

Harvey's legal embedding was trained on 20B tokens of US case law. If you're using it for Singapore, EU, or Croatian law — the semantic map may be miscalibrated for your jurisdiction. Check what your tool was trained on.

03

Invisible Architecture, Visible Risk

You never see which documents the AI retrieved. The embedding decision is made before you read a single word of the response. This is precisely why disclosure and verification obligations matter — Art. 47 CFREU and professional ethics alike.

The embedding model is the AI's mental map. A wrong map gives you a confident journey to the wrong destination.

From Retrieved Documents to Generated Answer

What the LLM actually does — and why vectors never really leave the room.


The full pipeline: embeddings → context window → attention → generation

Step 1 — Assembling the Context Window

The retrieved documents don't go into the LLM one by one. They are assembled into a single structured block of text — the context window — everything the LLM can "see" at once.

① System Prompt	Role, behaviour rules, output format instructions
② Retrieved Chunks	Top-20 document passages selected by the embedding model — cases, statutes, treatise sections — with citation metadata
③ Your Query	The actual question you asked, placed after the retrieved material
④ Conversation History	Prior exchanges in a multi-turn session — so the model knows what was already asked and answered

Key point: the LLM processes all of this simultaneously — not sequentially like a human reader.



transformer adds bits of context

store new info by modifying the word vectors in ways that are difficult for humans to interpret - new vectors, known as a hidden state, are passed to the next transformer in the stack

first few layers focus on understanding the syntax of the sentence – subsequent layers work to develop a high-level understanding of the passage as a whole

96 layers - word vectors with 12,288 dimensions—that is, each word is represented by a list of 12,288 numbers

57th layer - a vector for John would have encoded as a list of 12,288 numbers corresponding to the word John
contextual information: main character, male, married to Cheryl, cousin of Donald, from Minnesota, currently in Boise, trying to find his missing wallet

The final goal: all of the information necessary to predict the next word

Step 2 — Self-Attention: Every Word Talks to Every Other Word

The Q / K / V Mechanism



Query vector

"What am I looking for?"

This token searches for other tokens that help explain it.



Key vector

"What do I offer?"

This token advertises what it contributes to others.



Value vector

"Here's my content."

The actual information this token contributes when selected.

What This Means in Practice

Every token in the context simultaneously asks:

"Which other tokens matter most for understanding me?"

$Q \times K \rightarrow$ attention weight (how much to attend)

Weighted $V \rightarrow$ updated meaning of this token

Example:

"consideration" in a contract context attends strongly to "offer", "acceptance", "enforceable", "parties" — and weakly to everything else.

This is NOT programmed. It is learned from millions of legal texts during fine-tuning.

So What?

Hallucination Is Architectural, Not a Bug

01

The generation step samples from a probability distribution — it cannot distinguish 'I retrieved the right documents' from 'I retrieved nothing relevant.' When retrieval fails, generation continues anyway, producing fluent, confident, fabricated law. No version of this architecture eliminates that risk.

The Retrievable Is Auditable — The Reasoning Is Not

02

The embedding layer's decisions (which documents were retrieved, why) are, in principle, auditable — you can ask. The transformer's internal vector computations across 96 layers are not. For Art. 47 CFREU — the right to understand and contest — focus your challenge on the retrieval stage, not the generation stage.

A Jurisdiction Error at Level 1 Compounds Through Levels 2–5

03

If the embedding model (trained on US law) retrieves US cases for your EU/Croatian query, the LLM reasons correctly about those US cases — producing a coherent, beautifully-argued, wrong-jurisdiction answer. There is no self-correction mechanism. The error enters invisibly and exits confidently.

Confidence in the output is not evidence that the retrieval was correct.

Recital 61 – EU AI Act

- Certain AI systems intended for the administration of justice and democratic processes should be classified as high-risk, considering their potentially significant impact on democracy, **the rule of law**, individual freedoms as well as the right to an **effective remedy** and to a **fair trial**
 - AIs assisting judicial authorities in researching and interpreting facts and the law and in applying the law to a concrete set of facts.
- the risks of potential biases, errors and opacity
- the final decision-making must remain a human-driven activity
- does not extend to AI systems intended for purely ancillary administrative activities
 - anonymisation or pseudonymisation of judicial decisions, documents or data,
 - communication between personnel,
 - administrative tasks

The EU AI Act: Judicial AI as High-Risk

Recital 61 & Annex III, Point 8(a)

HIGH-RISK CLASSIFICATION

AI systems assisting judicial authorities in:

- Researching and interpreting facts and law
- Applying law to concrete facts

Rationale: Significant impact on democracy, rule of law, individual freedoms, and the right to an effective remedy.



Recital 61 Principle

"AI tools can support the decision-making power of judges, but should not replace it. Final decision-making must remain human-driven."



The Unresolved Tension

The opacity of AI makes it impossible for the judge — or the litigant — to evaluate the actual weight AI outputs carry in deliberation. Is the support/replacement distinction legally sustainable?

Art. 14 AIA — Mandatory human oversight for high-risk AI: understand capabilities, detect automation bias, interpret outputs, override or disregard AI output when necessary.

AI Tools Employed in Legal Practice



**LEGAL
RESEARCH**



**LEGAL
ANALYTICS**



**DOCUMENT
CREATION**




**LITIGATION
ASSISTANCE**



**OFFICE
ASSISTANCE**

AI models facilitating legal office management

- support automation of routine tasks like
 - case/client intake,
 - document generation,
 - file classification,
 - time tracking tool captures hours spent on specific tasks, automatically log time spent on activities like document drafting and email correspondence
 - assist with scheduling, monitor and resolve calendar conflicts, automatically proposes optimal meeting times while considering deadlines and availability,
 - provide automated updates on case status
 - billing
 - legal chatbots and virtual assistants facilitate client interaction (for example, by providing answers to common questions).
 - office assistants facilitating tasks like creating initial drafts of emails or letters or translations



The Foundational Regulatory Divide

HIGH-RISK (Annex III pt 8a)

- AI intended to assist a judicial authority
- Researching & interpreting facts and law
- Applying law to a concrete set of facts
- Alternative dispute resolution with legal effects
- Triggers Articles 9–15 AIA compliance obligations

ANCILLARY (Non-High-Risk)

- Does not affect administration of justice in individual cases
- Anonymisation / pseudonymisation (Recital 61, explicit)
- Communication between personnel
- Administrative & scheduling tasks
- Art. 6(3) exemption: narrow procedural / preparatory task

⚠ The binary is unstable — functional impact, not formal label, governs (anti-lacuna: SCHUFA C-634/21)

Eight AI Use Cases: Risk & Rights Overview

Use Case	AIA Status	Primary Art. 47 Concern	Key Risk
Case Triage & Routing	Formally ancillary; functionally contested	Access to justice (Art. 47(1))	Gatekeeping without contestation
Document Classification	Generally ancillary	Access / procedural track	Misclassification → wrong procedure
Anonymisation	Expressly ancillary (Recital 61)	Post-judgment GDPR	Accuracy; data leakage
Scheduling	Ancillary	Reasonable time (Art. 47(2))	Systemic deprioritisation
Speech-to-Text	Ancillary	Right of defence / appeal	Accuracy of record
Machine Translation	Ancillary	Linguistic access to proceedings	Accuracy; sparse language coverage
Process Automation	Variable	GDPR Art. 22; individual decisions	Automation of criteria judgments
Semantic Search	Ancillary	Epistemic integrity	Automation bias; structural narrowing

Case Triage & Routing — The Most Contested Category

What It Does

Sorts cases by subject matter, urgency, or procedural complexity before a judge sees them. Determines which cases proceed to full hearing, flagged as frivolous, or routed to abbreviated procedures.

The Art. 47 Problem

Where triage operates as a de facto access gatekeeper, individuals may be denied a hearing through a process they cannot see or contest (Art. 47(1) CFREU — access to justice; Heylens C-222/86 — right to reasons).

Regulatory Gap

Formally 'preparatory task' under AIA Art. 6(3) → non-high-risk. Functionally may cross into Annex III pt 8(a) if it conditions access to adjudication. Anti-lacuna principle applies.

Risk: Algorithmic Case Allocation

A hidden threat to 'tribunal established by law'



The Classification Paradox

Case allocation may fall within Art. 6(3) AIA exceptions as an 'ancillary administrative task' — possibly excluded from high-risk classification.

Yet it touches the core of Art. 47(2) CFREU: the 'tribunal established by law' guarantee. It determines who decides your case.

Opacity → Manipulation

Opaque allocation logic — proprietary, ML-weighted, undocumented — prevents parties from verifying cases were assigned by pre-existing rules, not post-hoc selection.

SCHUFA Anti-Lacuna Logic

C-634/21 SCHUFA: legal protection must follow functional reality of where power is exercised, not its formal institutional label. Allocation power exercised algorithmically still requires remedy.

Automation bias (Art. 14(4)(b) AIA acknowledged but no binding metric): over-reliance on algorithmic outputs compromises independence of judicial reasoning.

Automated Triage & Procedural Fairness

“

AI tools preventing cases from coming on for a hearing through automated triage implicate procedural fairness in a fundamental way.

— UNESCO Guidelines for the Use of AI Systems in Courts and Tribunals

Triage that filters cases before a hearing engages the right to a fair trial at its core — the decision to be heard is itself adjudicative.

Equality of Arms & the Article 47 Gap

1

EQUALITY OF ARMS

The asymmetry of timing

AI triage operates invisibly — before the adversarial process can even begin. A party cannot contest a filter they never see. The imbalance is temporal, not merely informational: by the time arms could be raised, the case has already been screened out.

2

THE ARTICLE 47 GAP

C-222/86 *Heylens* demands transparency

If AI triage acts as a first-instance filter on case access, the logic of *Heylens* applies: a decision restricting access to a hearing must carry reasons and be open to effective review. Opaque triage leaves no point at which that review can attach.

When Classification Becomes Triage

WHAT IT DOES

Document classification systems categorise incoming filings — complaint types, procedural categories, subject matter — routing each to the correct division, judge, or procedure.

THE RISK

Misclassification carries procedural consequences: incorrect limitation periods, the wrong procedural track, loss of urgent treatment.

THE TIPPING POINT

When classifying a document as failing the “general repercussion threshold” — the VICTOR criteria — means the case receives no substantive hearing, classification has become triage.

VICTOR — Brazil's Supreme Federal Tribunal

AI as the gatekeeper to the highest court in Brazil

The Problem

Brazil's STF receives hundreds of thousands of appeals annually. Each needs a 'general repercussion' constitutional gateway. Impossible to process manually.

What Kind of AI?

- Supervised Machine Learning + NLP (Portuguese text)
- Trained on 500,000+ STF decisions
- Neural network text classifier
- Pattern-matching (not generative AI — no hallucination risk)

What VICTOR Does

- Identifies cases matching already-decided repercussion themes → expedited route
- Flags cases likely not meeting the constitutional threshold
- Routes cases to the appropriate rapporteur

The Fundamental Rights Question

If VICTOR routes your case incorrectly, you may never reach a judge. Does that violate the natural judge principle? Does the opacity of the classifier satisfy Art. 47 CFREU's effective remedy guarantee?

Real Deployment #1 — VICTOR (Brazil STF)

5

seconds
per case
(vs. 44 min manual)

What It Is: ML + Deep Neural Networks; developed by Brazil's Supreme Federal Tribunal (STF) with the University of Brasília (2018).

How It Works: (1) OCR — converts images to text; (2) identifies document boundaries; (3) separates and classifies procedural documents; (4) identifies relevant 'general repercussion' themes and suggests them to court officials for validation.

The Threshold: Brazil's extraordinary appeal requires demonstrating 'general repercussion' (constitutional significance beyond the parties). VICTOR screens ~10,000 cases/year.

Art. 47 Risk: Automation bias — volume incentivises officials to accept the suggestion rather than perform independent review. No publicly available data on updates since initial deployment. Double black box: technical opacity compounded by informational opacity.

Real Deployment #2 — OLGA (Stuttgart OLG, Germany)

50%

projected
reduction in
processing time

What It Is: IBM NLU/NLP system; developed by IBM for the Stuttgart Higher Regional Court (Oberlandesgericht) under the Baden-Württemberg Ministry of Justice. Acronym: Oberlandesgerichtsassistent.


Context: Diesel emissions (Dieselgate) litigation — hundreds of thousands of near-identical consumer claims against car manufacturers, differentiated only by vehicle, engine type, mileage, and purchase date. Backlog >10,000 cases.

How It Works: Extracts structured metadata from unstructured legal text; classifies cases into judge-defined groups (vehicle model, engine type, mileage, year); surfaces relevant prior decisions; preserves full case history.

Design Strength: Grouping criteria are judge-determined — legally relevant variables explicitly specified by judicial experts, not inferred from training data. Significantly reduces opacity risk.

Residual Risk: If grouping criteria miss a legally relevant individual variation, the claimant is processed under a precedent that does not govern their case — Art. 47(1) access concern.

Inside the Classifier: Tokenise → Process → Classify



The classify step is where NLP quietly becomes triage.

Real Deployment #3 — iCTS (Singapore Courts)

What & How

Intelligent Court Transcription System

Developed by Singapore Courts in collaboration with A*STAR's Institute for Infocomm Research (I²R).

Technology: Domain-specific Automatic Speech Recognition (ASR) neural networks trained on legal vocabulary, proper names, and procedural terminology.

Function: Real-time transcription of court hearings — replacing manual transcription with an immediate, searchable verbatim record.

Analysis

Singapore philosophy: AI as augmentative, not adjudicatory. iCTS does not touch the substance of proceedings.

Transparency advantage: Indigenously built — the judiciary has access to technical specifications and training data. Avoids commercial black-box procurement.

Residual risk (Art. 47): Where a mistranscription becomes the authoritative hearing record, it directly affects the right to challenge witness evidence on appeal (see ECtHR, Murtazaliyeva v Russia, 2018 — adequacy of procedural records for effective appeal).

Real Deployment #4 — ANOPPI (Finland Ministry of Justice)

89.5%

name identification
accuracy
(pilot testing)

What It Is: AI pseudonymisation tool for Finnish court decisions. Project led by Finland's Ministry of Justice (initiated 2018; pilot 2020). Developed at the University of Helsinki. Open-source code.

How It Works: Combines (1) statistical NER — ML-based identification of probable personal data; and (2) rule-based morphological analysis — essential for Finnish as an agglutinative language with extensive inflectional forms (e.g. 'Matti Virtanen' → 'Matille Virtaselle' dative; 'Matti Virtasen' genitive; 'Mattia Virtasta' partitive). Includes coreference resolution — 'the plaintiff', 'Mr. V.', and 'the defendant in lower court proceedings' all identified as the same person.

Regulatory Status: Expressly ancillary — Recital 61 AIA explicitly lists pseudonymisation as non-high-risk. GDPR Arts. 5 & 25 (data protection by design) apply. Paradigm case of AI as GDPR compliance infrastructure.

EU Comparative Reach: Presented at EU Council Presidency seminar (2019). Croatia, Austria, Luxembourg now use automated anonymisation. Croatia adopted ANON software in early 2025. An EU model for GDPR-compliant judgment publication.

The 'Ancillary' Category as Regulatory Blind Spot

The AIA ancillary exemption rests on the intuition that administrative systems do not affect the administration of justice in individual cases.

In practice, several categories exercise material influence through indirect pathways:

Triage

Shapes which cases receive hearings

Classification

Determines procedural track and applicable precedent

Process Automation

May apply individualised criteria without human judgment

Semantic Search

Structures the informational environment of judicial reasoning

Art. 47 CFREU as Constitutional Floor: regardless of AIA categorisation, if a system materially conditions access to judicial protection, the effective remedy guarantee applies (SCHUFA anti-lacuna logic)

THE AI TOOLBOX

What Every Young Lawyer Must Know About AI in the Courtroom

6 Task Categories · 3 Critical AI Systems · Your Constitutional Obligations

Based on EU AI Act · CFREU Art. 47 · CJEU Case Law

6 Ways AI Enters the Courtroom

01

**AI-Assisted
Legal Research**

ChatGPT · Harvey AI · MARIA

02

**Digital Image &
CCTV Analysis**

CNN pattern recognition

03

**Deepfake
Detection**

Neural net classifiers

04

**AI Judicial
Reasoning Support**

Prometea · Shenzhen System

05

**E-Discovery /
Document Review**

Hansken · Predictive Coding

06

**Automated
Procedural Drafting**

Template + ML hybrid systems

Each category carries distinct Article 47 CFREU risks. We examine all six.

AI Tools Enhancing Legal Research

Swift Review of Databases of Legal Documents

- *Accelerate searches through case law, statutes, and regulations*
- Narrow searches to exactly the types of documents needed for specific task (not those that merely mention the keyword)
- *Efficiently locate pertinent precedents and supporting resources*

AIs Enhancing Legal Analytics

Using legal databases as a benchmark they scan documents to

- *Extract essential provisions and highlight key arguments*
- *Create various summaries of legal arguments*
- *Identify relevant missing statutory provisions or court rulings*
- *Recommend additional pertinent cases or legal research*
- *Verify references and citations*
- *Recognize trends within applicable case law*
- Identify how contract language may deviate from other standard uses
- Determine what language particular parties tend to prefer
- Evaluate favourability of particular contract provision in light of the parties' usual behaviour

AI Tools Facilitating Creation of Legal Documents

- Scan vast data sets to
 - assist by generating (different semantic versions of) drafts of legal documents in compliance with jurisdiction-specific requirements
 - learn from sets of existing contract templates and drafting preferences to generate customized drafts of legal documents
 - filling in templates with case-specific details
 - create “clause library” that can be used to suggest appropriate language based on context and jurisdiction
 - identify how contract language may deviate from other standard uses
 - determine what language particular parties tend to prefer
 - evaluate favourability of particular contract provision in light of the parties usual behaviour

Two Fundamentally Different AI Worlds

Before we explore each tool — you need this distinction in your head permanently.

DETERMINISTIC / SUPERVISED ML

- ✓ Trained on curated legal data
- ✓ Produces traceable outputs
- ✓ Consistent on similar inputs
- ✓ Auditable error modes

Examples:

Prometea · MARIA · VICTOR · Hansken

GENERATIVE AI / LLMs

- ⚠ Trained on vast unstructured data
- ⚠ Generates text probabilistically
- ⚠ Cannot verify its own outputs
- ⚠ Confabulation risk = INVISIBLE errors

Examples:

ChatGPT · Harvey AI · Copilot

This distinction determines constitutional risk. A wrong citation from an LLM is invisible — it looks identical to a real one.

AIs Providing Litigation Support

E-Discovery: sorting through massive volumes of electronic data during litigation

Provide precise, contextually relevant answers to concrete legal prompts

Find specific cases, statutes, and precedents

“Trial assistants” tools generate relevant case law, provide guidance on rules of procedures, suggest possible counter-arguments and objections if prompted

Organize case materials in chronological order and points out where evidence might be missing

Process multiple file types, from emails to social media posts, identifying relevant evidence and potential exhibits

Examine witness statements and deposition transcripts to identify patterns that might indicate credibility issues, identify contradictions and inconsistencies across multiple testimonies

Creates visual representations of case events, making it easier to present complex narratives to judges

Tasks 02 & 03: Visual Evidence in the Age of AI

CCTV & IMAGE ANALYSIS

The Golden Rule:

Enhancement = showing more clearly what is already there.
NOT inventing new detail.

AI may legally:

- Improve clarity of existing features
- Recognize license plates / faces
- Cross-reference footage across feeds

AI may NOT:

- 'Fill in' detail it predicts should be there
- Generate a 'nicer' version of the image
- Output without a full audit trail

DEEPPAKE DETECTION

What is a deepfake?

Synthetically generated audio/video. Can be the OBJECT of proceedings (harassment) or fake EVIDENCE in other proceedings.

UK Custody Case:

Deepfake audio submitted to portray a father as violent. If undetected: catastrophic for his rights.

The Arms Race Problem:

Detection accuracy < 70% even for best systems (Meta Challenge, 2,114 teams). And it keeps getting harder.

Your Competence Obligation

ABA Rule 1.1 / Professional conduct standards across EU jurisdictions now include technology literacy.

 Know what type of AI produced the output you are relying on

 Verify every citation generated by an LLM against primary sources

 Demand audit trails for any AI used in evidence processing

 Challenge AI evidence you cannot scrutinize — Art. 47 gives you that right

The Colombian Constitutional Court (T-323/2024): AI must support — never replace — human judgment.

LAW × ARTIFICIAL INTELLIGENCE

AI in the Courtroom

Decision-Making, Judgment & Sentencing

Predictive Justice

Sentencing Algorithms

Appeal Screening

Automated Scoring

AI in ADR

Predictive Justice

Can an algorithm read the court's mind?

What It Does

- Takes the facts of a case as input
- Matches against corpus of past decisions
- Outputs: probability of outcome, relevant precedents
- Used by: law firms (Case Law Analytics) OR courts (Prometea)

PROMETEA (Argentina): predicts outcome in

20 seconds · 96% accuracy

⚠ The Legal Problem

Self-fulfilling prophecy

If judges anchor to the prediction, AI stops predicting outcomes and starts determining them.

Historical bias

Trained on past case law. Past injustice → future injustice, at scale and speed.

Equality of arms

Only well-resourced parties (state, corporations) can afford these tools. Art. 47 CFREU problem.

SÓCRATES — Brazil's Superior Court of Justice

When AI not only screens cases — but drafts the decision

01

Case Batching

Groups legally identical appeals. One precedent resolves thousands. Massive efficiency — but also massive homogenisation.

02

Legal Research

Retrieves relevant STJ precedents and arguments based on the case's key legal questions. Embedded AI legal assistant.

03

Admissibility Pre-Screening

Assists in identifying which cases raise genuine federal law questions — the STJ's gateway test.

04

Draft Text Generation

Generates initial drafts of procedural decisions for repetitive cases. Human judge 'signs off.'

AUTOMATION BIAS RISK: When AI drafts the decision, is the judge deliberating — or rubber-stamping? SCHUFA says nominal human sign-off is not enough.

PretorlA — Colombia's Constitutional Court

AI in the service of fundamental rights — and its paradox

What PretorlA Does

Triage & Routing

Classifies tutela petitions by constitutional right category

Urgency Flagging

Detects imminent threats to life, health crises, children — flags for emergency processing

Repetitive Claim Grouping

Identifies thousands of identical violations (e.g. same insurer denying same medication) → one structural ruling with erga omnes effect

Precedent Research

Surfaces relevant Constitutional Court jurisprudence for judicial staff

The Paradox

PretorlA is designed to protect constitutional rights. Yet the citizens it classifies — often the most vulnerable — have no visibility into how the classification was made or how to challenge it.

The AI meant to defend fundamental rights may itself create an Art. 47 CFREU deficit.

T-323/2024 — Colombian Constitutional Court

A judge used ChatGPT in a tutela ruling. Court held: AI may support managerial, documentary tasks — but NOT substitute human judgment in the actual decision. AI-generated text ≠ judicial reasoning.

COMPAS and the 'Robo-Judge' threshold

COMPAS — Correctional Offender Management Profiling for Alternative Sanctions

Developed by: Equivant (formerly Northpointe), USA

Function: Scores offender recidivism risk (1–10 scale)

Used by: US courts at sentencing (via function creep — originally a correctional placement tool)

Key case: State v Loomis (Wisconsin SC, 2016)

ProPublica (2016): Black non-reoffenders were

2× more likely to be wrongly flagged as high-risk

The Spectrum

AI ASSIST

Informs one variable in sentencing

AI INPUT

Score presented to judge as key factor

ROBO-JUDGE

Fully automated sentence — NO human

The Bias Insufficiency

- Four major concerns
 - algorithm design and data flaws
 - cognitive bias
 - prejudice
 - misuse
- difficult to detect; harmful to the fairness and legitimacy of legal systems
- AI trained on problematic data will generate problematic outputs
 - Data is not neutral, it has history; generated solution will mirror history, perpetuate its structure
 - perverse feedback loops
 - the tool continue to learn and improve from growing dataset which it expends through optimizing its own functioning that had incorporated prejudice and biases built in original data

Risk III: The 'Black Box' Paradox

Inscrutability and the duty to reason

TECHNICAL BLACK BOX

Millions of weighted parameters in convoluted neural networks. The transformation from input to output cannot be rendered understandable — by design.



LEGAL BLACK BOX

Trade secret protections (Art. 78 AIA, GDPR Art. 15(4)) create a corporate secrecy layer. Formal compliance ≠ functional intelligibility.

COMBINED EFFECT: Structural impossibility of effective contestation — directly undermining the adversarial principle of Art. 47 CFREU

Opacity Variations

Intentional, arising from corporate or state secrecy;

Technical, due to the complex nature of AI;

Intrinsic, stemming from the fundamental differences in human and algorithmic cognition.

The CJEU Framework: Key Rulings

C-817/19 Ligue des droits humains (2022)

Holding: Parties must be able to understand how AI criteria and programs work. Opacity may deprive subjects of effective judicial remedy.

Implication for judicial AI: *If even administrative authorities must ensure individual human review of every AI output, a fortiori a judge relying on AI must demonstrate meaningful deliberation — not formal endorsement.*

C-634/21 SCHUFA Holding (2023)

Holding: Legal protection must follow functional reality of where power is exercised — not its formal label. An AI score that plays a determining role is the operative decision.

Anti-lacuna logic: *Where algorithmic power drives an outcome but is labelled 'preparatory', Art. 22 GDPR and Art. 47 CFREU close that gap. Human sign-off as formality is not sufficient.*

Principle: Genuine judicial deliberation, not algorithmic endorsement, is what Art. 47 CFREU requires.

Compas only points the direction

- COMPAS example:
 - satisfies „equal positive” predictive values but fails „equal false positive” predictive values
 - designers provided data indicating algorithm fairness since the number of defendants classified as ‘high risk’ that reoffended was the same whether they were black or white (around 60 per cent of high-risk defendants).
 - NGO ProPublica analysed the same data to show that among defendants who ultimately did not reoffend, blacks were more than twice as whites to be classified as likely to reoffend
 - satisfying both fairness criteria is difficult – assuming the existence of racial disparity in recidivism rates the model is highly likely to produce either false positives or false negative
 - Which criteria of fairness is proper? Equalised odds (equal balance in error/misclassification rates) or predictive parity (predict risks at similar rates across different groups)?

The Double Black Box • The Anti-Lacuna Principle • Your Role

THE DOUBLE BLACK BOX

Technical opacity (you can't inspect the algorithm) + Legal opacity (trade secrets, Art. 78 AIA, GDPR Art. 15(4)) operate together. The result: citizens and lawyers cannot see, understand, or effectively challenge the AI that determines their rights.

1

The Anti-Lacuna Principle

SCHUFA: courts will look past formal structure to functional power. If AI determines the outcome, it IS the decision — no matter how many humans 'sign off.'

2

Human in the Loop ≠ Human in Control

Volume + automation bias means nominal human review is often fictional. Your client has a right to genuine deliberation, not rubber-stamping.

3

Your Obligation as a Lawyer

Ask: What AI was used? Who trained it? On what data? Can it be challenged? Request the Art. 15(1)(h) GDPR explanation. Invoke Art. 47 CFREU. The law gives you these tools.

Safeguards for Impartiality in Digital Justice Systems



Legal Safeguards

- › Art. 14 AIA: Mandatory, genuine human oversight — not rubber-stamping
- › Art. 86 AIA: Transparency & explanation obligations with anti-discrimination extension
- › Art. 47 CFREU: Systematic auditing, mandatory disclosure of algorithmic logic
- › Ligue / SCHUFA doctrines: Remedy follows functional exercise of power



Institutional Safeguards

- › Structural quality control — inquisitorial oversight of AI outputs
- › In-house training on curated official court documents only
- › Data protection: prohibition on public AI tools processing confidential proceedings
- › Allocation systems: documented, auditable, challengeable rules



Ethical & Cultural Safeguards

- › Ethical standards as complement to legal rules — not substitute
- › Institutional culture of critical engagement with AI outputs
- › Binding metrics for distinguishing deliberation from algorithmic endorsement
- › Judicial training: recognising automation bias in practice

In-house trained AI on curated court datasets should be considered a baseline requirement — addressing both narrowness risks and confidentiality obligations simultaneously.

