Foundation of Data Protection Professionals in India
[Not for Profit, Section 8 Company limited by guarantees]
CIN No: U72501KA2018NPL116325: GSTN 29AADCF4963H1ZC
Registered Office: No 37, "Ujvala", 20th Main, BSK first Stage, Second Block, Bangalore 560050
Web: www.fdppi.co.in : www.fdppi.in : E Mail fdppi@fdppi.in: Ph: 08026603490: Mob: +91 8310314516

Date: 19 June 2026

To

The Member Secretary,

AI Committee, Supreme Court of India

(office.regcc@sci.nic.in)

Ref: Notice dated 03.06.2026 inviting views/suggestions of stakeholders and the general public (last date: 20.06.2026)

Dear Sir

We are pleased to submit our views on the proposed regulations for AI usage in Judiciary. This submission is being made on behalf of Foundation of Data Protection Professionals in India (FDPPI)

FDPPI is a not-for-profit Section 8 company promoted by individual data protection professionals not aligned to any Big Tech company nor NASSCOM and therefore does not carry any vested interests in AI vendors.

The undersigned, known popularly as Naavi, is a pioneer in the field of Cyber Law and Data Protection, founder of www.naavi.org, Founder Chairman of FDPPI and Chair Professor of FDPPI's AI Chair. He is the architect of the Data Governance and Protection Standard of India (DGPSI) and its artificial-intelligence extension, DGPSI-AI. This is a self-regulatory standard for AI deployment by Data Fiduciaries for DPDPA compliance, and a prospective regulatory seed.

FDPPI has also developed an "Independent Data Auditor" ecosystem and is establishing a Centre of Excellence in "AI Risk Management" in association with an academic institution both directly relevant in the present context.

General observations

We respectfully commend the Committee for the formulation of these AI implementation guidelines though it is noted that it is applicable to the Judicial sector. The draft is among the

most comprehensive judicial AI governance frameworks proposed anywhere in the world and is, in substance, ready to be taken forward as a formal instrument.

We particularly welcome the principle of human primacy in Section 4, which provides that the use of AI shall at all times remain strictly subservient to human judgement and judicial authority. In our view this is the single most important feature of the framework, and it is fully consistent with the foundational principle of DGPSI-AI, that behind every AI algorithm there shall be an identifiable human responsible for its outcomes. The absolute and non-derogable prohibitions in Regulation 20 in particular the bar on algorithmic adjudication, risk scoring and behavioural profiling likewise deserve special appreciation and should be retained without dilution.

Because this framework is issued under the authority of the Hon'ble Supreme Court, it will inevitably serve as a due-diligence benchmark for legal practitioners, government agencies and private-sector AI developers operating in India. We therefore urge that the final instrument be drafted with that wider influence in mind.

Our suggestions below, offered in good faith, are summarised in this Memorandum; the detailed regulation-by-regulation comments appear in Annexure A, and a focused note on the definition of "AI" appears in Annexure B.

Governance structure — cost, scale and the risk of over-administration

While the importance the draft attaches to AI governance is entirely appropriate, the proposed architecture an Apex Body ("Appropriate Authority"), five sub-committees, the Centre of Research and Excellence on Artificial Intelligence (CoRE-AI), and AI Committees with supporting AI Secretariats at the Supreme Court and at every High Court may in our humble opinion create a disproportionately large administrative establishment, with attendant cost and duplication and can be refined.

We see real scope for reducing this cost through a simplified organisational design, in particular by drawing on officers deputed from other Government organisations or academic institutions.

We would respectfully suggest that the structure be revisited, before finalisation, by a small independent review group for example drawing on the Indian Institute of Science, an IIM, an IIT, a National Law University, an organisational-design expert from the private sector, and a senior Chartered Accountant. Such a review would help pre-empt concerns that the Central Government and the Comptroller and Auditor General may otherwise raise.

Centralisation of oversight and approval functions

Our most significant structural recommendation concerns the proposed decentralised structure of the oversight. The draft contemplates that approvals, technical and ethical impact assessments, incident reporting, content verification, business-continuity arrangements and review of legacy systems will be handled at the level of individual AI Committees, including at each High Court. This risks duplication of effort, conflicting decisions between Courts, and substantial recurring cost.

To preserve the independence of the High Courts while securing coordination, we recommend creation of an AI General Council under the Chairmanship of the CJI and comprising the Chief Justices of all High Courts. This could be the apex body driving this regulation. Each Chief Justice

of the High Court may be supported by a personal secretariat of not more than two persons, essentially for coordination.

The following committees should be centralised under the Apex Body / a single principal AI Secretariat:

- Software approval (Chapter III): a single, centralised Technical Committee, so the same product is not separately evaluated by different Courts;
- A central violation and de-listing register, so a non-compliant tool or vendor can be removed across all Courts at once;
- The format and conduct of the Technical and Ethical Impact Assessment (R.35), the business-continuity / fall-back process (R.42), review of legacy systems (R.41), cyber-security (R.15) and content monitoring (R.44).

Technical evaluation, vendor assessment and standard-setting may thus be centrally coordinated instead of each High Court having a local control. Since the entire scheme works under the Central AI General Council, the autonomy of the High Courts is preserved.

Data protection — status of Courts under the DPDPA, 2023

While addressing the data protection requirements, the draft introduces the term "Sensitive Judicial Data" which is applicable to personal data of parties and witnesses. The DPDPA does not define "sensitive" personal data; it defines only the "Significant Data Fiduciary."

To maintain the harmony between the regulations and DPDPA, Judicial systems may either be treated as "Significant Data Fiduciaries" or entities exempted as Instrumentalities of state engaged in maintenance of public order.

We therefore recommend that the status of judicial institutions under the DPDPA may be clarified by MeitY through an appropriate notification or interpretative guidance.

Further, as the draft specifies no internal appeal mechanism, the grievance redressal (related to Data Principal Rights has to proceed through the Data Protection Board (unless exempted). This may be stated expressly.

Audit philosophy — independence of assurance

The draft favours an in-house audit model and restricts the sharing of source code, algorithms and datasets with third parties. We appreciate the security rationale, but an audit conducted entirely in-house carries an inherent limitation of independence, this procedure of internal audit only may be considered as lacking in full objectivity.

We recommend a balanced model in which internal audits are supplemented by accredited, security-cleared independent AI and data auditors operating strictly under controlled conditions on Court premises, under confidentiality undertakings, with no source code, model weights or datasets leaving Court custody.

Such auditors may be accredited under a competency framework covering law, AI governance, cyber-security, privacy and evidence management.

Procurement and engagement of the private sector (Regulation 46)

We welcome the rigour of Regulation 46 ownership of data and outputs, purpose limitation, explainability, indemnity in favour of the Courts, on-premise/sovereign-cloud deployment for sensitive judicial data, the bar on retraining on Court data without approval, and retention by Courts of ownership or a perpetual royalty-free licence.

We however draw attention to the practical effect on legacy systems. Applied without transition, every existing vendor claiming any AI capability would need fresh clearance. The one-year compliance-review window is welcome and the transition arrangements should be made explicit.

Frameworks such as DGPSI-AI may assist vendors in demonstrating readiness; vendors may be encouraged first to satisfy the recommended "AI Developer" implementation specifications, which cover all the requirements envisaged in the regulation and more.

The AI Content Verification Authority (Regulation 44)

Regulation 44 appears to be a little ambitious. As drafted, it could imply a forensic facility tasked with verifying every item of AI-generated content used in court processes. Current AI-detection technologies are themselves probabilistic and may not provide conclusive evidence of AI generation. We suggest its scope be clarified and, to the extent retained, that it be operated centrally rather than replicated at each Court.

"Innovation over Restraint" — toward "Restrained Innovation" (Regulations 16–17)

The draft commendably states a presumption in favour of responsible adoption. In substance, however, the detailed provisions are weighted towards restraint and are in places stringent.

We suggest the Committee to take steps to ensure that the operative clauses give genuine effect to the stated preference for responsible innovation, so that beneficial, well-governed tools are not discouraged, while equally ensuring the regulation does not yield to the influence of a commercially powerful AI-developer world.

FDPPI suggests the guiding principle of **"Restrained Innovation"** encouraging innovation while maintaining judicial reliability and public trust.

Capacity building (Chapter VIII)

We strongly support the structured training mandate in Regulation 49 and the "living repository" of best practices in Regulation 51..

Miscellaneous — evidentiary standards and risk classification

1. As AI-generated content increasingly forms part of judicial records, guidance may be developed on disclosures in certificates under Section 63 of the Bharatiya Sakshya Adhiniyam wherever AI-assisted evidence generation, analysis or presentation is involved.

   A risk classification of AI systems and associated data, Low, Medium, High and Critical may be explored and applied to both personal and non-personal data, for example:

   Low Risk: translation, transcription, scheduling;
   Medium Risk: legal research, summarisation;
   High Risk: predictive analytics, risk scoring;
   Critical Risk: any AI affecting adjudicatory outcomes.

2. Any AI failure that materially affects liberty, access to justice, fair hearing, or procedural fairness should be treated as a Judicial Safety Event and subjected to mandatory incident reporting

Consolidated summary of major recommendations

1. Constitute an AI General Council of all Chief Justices of High Courts, chaired by the CJI, as the apex policy-making body for the entire ecosystem.
2. Support each High Court Chief Justice with a personal secretariat of not more than two persons for AI coordination.
3. Unify all other functions — technical evaluation, vendor assessment, standard-setting, audit coordination, content monitoring — under central Technical and Governance Committees, and have the overall structure independently reviewed for cost and duplication.
4. Request MeitY to clarify the status of judicial institutions under the DPDPA, 2023 (including Section 17(2)), through notification or interpretative guidance.
5. Supplement internal audit with external audit by independent, security-cleared AI and data auditors under controlled, on-premise conditions.
6. Make explainability an authenticated vendor deliverable for all AI Systems and fix a named human accountable (Designated Officer and vendor Human Handler) for every system in the AI Register.
7. Introduce a graded risk classification (Low/Medium/High/Critical) for both personal and non-personal data, with Critical-Risk systems requiring Apex Body approval, behaviour monitoring and tamper-proof kill switches.
8. Adopt a corrected, OECD-aligned definition of “AI” as the regulatory gate, layered with the risk classification above and a vendor-confirmed functional test (see Annexure B), and adopt the principle of “Restrained Innovation” with guidance on Section 63 of the Bharatiya Sakshya Adhiniyam for AI-assisted evidence.

9. Any AI failure that materially affects liberty, access to justice, fair hearing, or procedural fairness should be treated as a Judicial Safety Event and subjected to mandatory incident reporting

India has the opportunity to create the world's first comprehensive judicial AI governance framework balancing innovation, accountability and constitutional values. FDPPI believes that with modest refinements to governance architecture, audit independence, evidentiary standards and risk-based implementation, the draft can become a model framework for the responsible judicial adoption of AI.

FDPPI would be privileged to engage further including through a detailed clause-by-clause note and participation in any consultation the Committee may convene — and to place the resources of its professional community at the Committee's disposal.

With respect and regards,

Vijayashankar Na (Naavi)
Founder Chairman
& Chair Professor, FDPPI AI Chair

Annexure A — Section-by-Section Comments

The following table sets out FDPPI's detailed comments mapped to the relevant provisions of the draft Regulations, with the rationale for each and, where applicable, the corresponding DGPSI-AI reference (Principle / Deployer Model Implementation Specification (MIS) / Developer (Dev.) MIS).

Regulation / Provision	Comment & Suggestion	Rationale	DGPSI-AI Ref.
Reg. 3(1)(m) & 3(1) [new]	Adopt a functional, vendor-confirmed test for what constitutes "AI" — i.e. whether the software leverages autonomous learning or probabilistic models to adapt behaviour and generate outputs not fully predetermined by explicit code. Insert a definition of "AI Agent" / "agentic AI." (See Annexure B for the comparison with the DGPSI definition, and Annexure B.7 for proposed corrected definition text.)	Makes the boundary operational for procurement and approvals and prevents disputes over whether a tool falls within the Regulations. Reg. 23(a) already refers to "autonomous AI agents" but leaves the term undefined despite its higher risk profile.	Deployer MIS-1
Regs. 16–17	Counterbalance "Innovation over Restraint" with a precautionary corollary — where the risk of an AI System cannot reasonably be estimated, it is presumed significant and attracts the safeguards of Reg. 12(2) until experience justifies reclassification. Where an authority treats an unknown risk as not significant, it records an "AI Deviation Justification Document." FDPPI commends the principle of "Restrained Innovation."	AI risks (emergent behaviour, hallucination, model drift, adversarial manipulation) are unquantifiable at deployment. The presumption of adoption should never operate in a risk vacuum, nor should the framework yield to a commercially powerful AI-developer lobby.	Principle 1; Deployer MIS-3
Reg. 7 r/w 46(4)(h)	Make explainability an authenticated deliverable signed by the AI Service Provider, describing algorithmic functioning, data inputs and foreseeable harms. Require it for all AI Systems, not only "High-Risk" tools (a term not defined in Reg. 3). Accompany party disclosures under Reg. 43(1) with an accessible explainability summary.	Authentication converts a marketing document into an accountable representation on which the Court can rely; undefined "High-Risk AI Tool" should be defined or replaced.	Principle 3; Dev. MIS-1, MIS-6
Reg. 8 r/w 3(1)(u) & 37	Make nomination of a Designated Officer mandatory for every approved AI System, recorded in the AI Register. Require the licensing contract to name the vendor's own "Human Handler" with business contact details.	Creates an unbroken human chain of accountability from developer to deployer to user.	Principle 2; Deployer MIS-4, MIS-5; Dev. MIS-2
Reg. 12 & 24(1)(a)	Task the Apex Body with notifying a graded risk classification with a top "Critical Risk" tier (high autonomy + sensitive judicial data + decision-support). Critical Risk systems to require Apex Body approval, continuous behaviour monitoring in the AI Register and shorter audit cycles. Suggested tiers: Low (translation, transcription, scheduling); Medium (legal research, summarisation); High (predictive	Gives concrete content to the proportionality principle and channels the most dangerous systems to the highest level of scrutiny.	Deployer MIS-9; Dev. MIS-12

Regulation / Provision	Comment & Suggestion	Rationale	DGPSI-AI Ref.
	analytics, risk scoring); Critical (any AI affecting adjudicatory outcomes). Apply to both personal and non-personal data.		
Reg. 15 / 48(5)	Manage cyber-security and cyber-security audits as a centrally coordinated function supported by specialist capability, rather than replicated at every Court.	Specialist security capacity is scarce; central coordination avoids duplication and uneven protection.	Deployer MIS-2
Reg. 19(1)(f) & 35(3)	Condition approval of litigant-facing conversational AI on a documented guardrail statement: (i) no manipulative design/nudging; (ii) clear disclosure that the user is interacting with AI, not Court staff; (iii) prominent disclaimer that outputs are not legal advice or orders; (iv) escalation to a human officer. The Ethical Impact Assessment to evaluate these guardrails.	Unrepresented litigants interact with chatbots most directly and are most exposed to dark patterns and manipulation.	Principle 5
Reg. 23(a)	Require the Apex Body to issue dedicated standards for agentic AI within a defined period: scope-of-action limits, mandatory logging of every autonomous action, Human-in-the-Loop checkpoints for any action with legal effect, classification of fully autonomous agents as Critical Risk by default, and post-deployment behaviour monitoring.	Agentic systems capable of multi-step autonomous action present a materially higher risk than static models and are entering knowledge-work products rapidly.	Dev. MIS-12
Reg. 35	Where an AI System processes personal data, clarify that a TEIA (Technical and Ethical Impact Assessment) in the prescribed format shall also serve as the DPIA under the DPDPA. Add an "AI Justification Document" recording why an AI-led process is preferred over a non-AI alternative (technical, operational, economic).	Avoids duplicative assessments, dovetails with the DPDPA regime referenced in Reg. 47, and creates a reviewable record for audits.	Principle 4; Deployer MIS-1, MIS-7
Reg. 38(2)	Replace the absolute in-house-only audit rule with audits by independent auditors empanelled and security-cleared by the Apex Body, performed strictly on Court premises (or within the Controlled Environment) under binding confidentiality undertakings, with no source code, model weights or datasets leaving Court custody. Accredit such auditors under a competency framework covering law, AI governance, cyber-security, privacy and evidence management.	In-house teams audit systems their own institution approved — a structural conflict of interest. Independence and scarce technical depth (bias testing, adversarial robustness, drift) are the foundation of audit credibility. Sits uneasily with the audit rights in Regs. 46(4)(e) and 48(5).	Deployer MIS-2; Dev. MIS-11
Reg. 41	Conduct the review of legacy/existing systems as a centrally coordinated exercise; make the one-year transition arrangements explicit.	Avoids inconsistent treatment of the same product across Courts and clarifies vendor obligations.	Deployer MIS-1

Regulation / Provision	Comment & Suggestion	Rationale	DGPSI-AI Ref.
Reg. 42 r/w 46(4)	Add as mandatory contract provisions: a tamper-proof kill switch controlled independently of the model (which cannot access or override its own termination); documented vendor emergency-handling instructions feeding the fall-back protocol and training curriculum; verification of these during Controlled Environment Testing for Critical Risk systems. Operate the fall-back/business-continuity process centrally.	Reg. 42 provides an institutional fall-back but not the engineering capability to halt a malfunctioning system itself.	Deployer MIS-9; Dev. MIS-8–MIS-10
Reg. 44	Clarify the scope of the AI Content Verification Authority and, to the extent retained, operate it centrally rather than at each Court. Recognise that current AI-detection technologies are themselves probabilistic and may not conclusively establish AI generation.	As drafted it could imply a forensic facility verifying every item of AI content — a potentially unfunded mandate resting on unreliable detection.	—
Reg. 46(4)	Add to the mandatory contract provisions: (i) assurance, preferably third-party tested, that the software is vulnerability-tested, secured for confidentiality/integrity/availability of Court data and free from malware; (ii) documentation of the training and testing process; (iii) documentation of default configuration and any re-configuration/re-training in normal use; (iv) vendor disclosure of its own use of AI agents in developing/maintaining the system. Extend the indemnity under 46(4)(i) to litigants and third parties harmed by defective vendor systems. Vendors may demonstrate readiness via the DGPSI-AI "AI Developer" specifications.	Strengthens assurance at the points outside the Court's control (vendor, model, supply chain) and ensures the affected citizen is not left to the general law alone under Reg. 53.	Principle 4; Deployer MIS-8; Dev. MIS-3, 6, 7, 13
Reg. 48(4) & 27	Adopt a "fading memory" time-weighting of learning data — a time-sensitivity parameter tied to the age of the observation — as a technical standard specified by the Technical Committee.	Prevents superseded precedents, amended statutes and outdated practice from continuing to drive model behaviour, especially in legal-research and summarisation tools (Reg. 19(1)(d)).	Deployer MIS-9(5)
Reg. 53	State the grievance and appeal pathway expressly; provide training for grievance-redressal personnel distinct from that of general AI users.	The draft specifies no internal appeal mechanism; read with Reg. 53 an aggrieved person may approach other competent fora, including the Data Protection Board route.	—
DPDPA 2023 – "Sensitive Judicial Data" / s.17(2)	Seek a MeitY notification or interpretative guidance clarifying the status of judicial institutions under the DPDPA. The DPDPA does not define "sensitive" personal data;	Avoids unintended enhanced-compliance obligations (DPO, DPIA, audits) and resolves	—

Regulation / Provision	Comment & Suggestion	Rationale	DGPSI-AI Ref.
	treating all judicial data as "Sensitive Judicial Data" may, by implication, make Courts Significant Data Fiduciaries. Complete exemption arises only under s.17(2); s.17(1) exemptions are partial and do not displace the s.8(5) security obligation.	perceived conflict between the framework and the DPDPA.	
Governance architecture (Ch. IX; Regs. 24–30)	Constitute an AI General Council of all Chief Justices of High Courts, chaired by the CJI, as the apex policy body, preserving the administrative independence of each Court. Each High Court CJ to be supported by a secretariat of not more than two persons for coordination. Unify technical evaluation, vendor assessment and standard-setting under central Technical and Governance Committees. Have the structure reviewed by an independent group (IISc, an IIM, an IIT, an NLU, a private-sector organisational-design expert and a senior Chartered Accountant).	The proposed multi-tier structure (Apex Body, five sub-committees, CoRE-AI, plus AI Committees and Secretariats at every High Court) risks a disproportionately large, costly establishment, duplication and conflicting decisions — concerns the Central Government and CAG may otherwise raise.	—
Bharatiya Sakshya Adhiniyam – s.63	Develop guidance on disclosures to be incorporated in certificates under s.63 of the Bharatiya Sakshya Adhiniyam wherever AI-assisted evidence generation, analysis or presentation is involved.	As AI-generated content increasingly forms part of judicial records, evidentiary integrity requires clear disclosure standards.	—

Abbreviations: TEIA — Technical and Ethical Impact Assessment; DPIA — Data Protection Impact Assessment; DPDPA — Digital Personal Data Protection Act, 2023; MIS — Model Implementation Specification; HITL — Human-in-the-Loop.

Annexure B — Comparative Note on the Definition of "Artificial Intelligence"

Purpose. This note compares the definition of "Artificial Intelligence" in the draft Regulations with the definition adopted under DGPSI, and recommends a combined approach for Regulation 3(1)(m). It supplements the first row of Annexure A.

B.1 The two definitions

Draft Regulations. *"a machine-based system that infers, learns, and generates decisions, predictions, and recommendations from data, with a varying degree of autonomy, such as, algorithms, computational processes, and software, deployed for court processes, excluding general-purpose software or digital tools, unless such software or tools are specifically embedded with, augmented by, or functionally dependent upon, artificial intelligence."*

DGPSI. *"AI is a class of automated data processing system where the human intervention in decision output and application of decision to a business decision is below an acceptable threshold."* DGPSI operationalises the threshold through three classes: **Class 1** — software with code-correcting ability without the intervention of a human developer; **Class 2** — a system that automatically implements a decision affecting a human; and **Class 3** — a system that reacts to human emotions or is capable of creative output, including generative AI.

B.2 Two different questions

The two definitions belong to different traditions. The Court's definition answers an ontological question — what kind of thing is an AI system — and answers it descriptively, by listing capabilities. The DGPSI definition answers a regulatory question — at what point must the safeguards that attach to AI come into play — and answers it by reference to the displacement of human control. The first draws a boundary around a category of technology; the second draws a boundary around a category of risk. Almost every other contrast flows from this distinction, summarised below.

Dimension	Supreme Court draft definition	DGPSI definition
Underlying question	What kind of thing is an AI system? (ontological / descriptive)	At what point must AI safeguards attach? (regulatory trigger / functional)
Defining axis	Capabilities of the system — it "infers, learns, and generates" with "varying autonomy."	Loss of human control — human intervention "below an acceptable threshold."
Lineage	Tracks the OECD definition and Article 3(1) of the EU AI Act — internationally aligned.	Indigenous, accountability-first; coheres with DGPSI-AI Principle 2 (one human behind every algorithm).

Dimension	Supreme Court draft definition	DGPSI definition
Scope	Domain-bound — limited to "court processes," with a general-purpose-software carve-out.	Domain-neutral — any "business decision"; portable across sectors.
Treatment of generative AI	Outputs listed as "decisions, predictions, and recommendations" — "content" not named; possible gap.	Class 3 expressly covers creative/generative and affective systems.
Self-learning systems	"learns" appears in the verb list but is not elaborated.	Class 1 targets self-correcting behaviour without a human developer — a concrete trigger.
Administrability as an approval gate	Hard at the margin — "functionally dependent upon AI" is difficult to certify cleanly.	Easier — a vendor can attest "this is a Class 2 system," but the threshold is relative.
Principal weakness	Circular (defines AI by reference to AI); conjunctive verbs ("infers, learns, and generates") read as requiring all three.	"Acceptable threshold" undefined; classes drift between a control axis and a capability axis and may overlap.

B.3 The draft Regulations' definition — assessment

The Court's wording is modelled on the OECD definition and Article 3(1) of the EU AI Act. That lineage is a strength: it aligns Indian judicial practice with the emerging global consensus, is defensible against the charge of idiosyncrasy, and eases future interoperability. The general-purpose carve-out is also sensible — it keeps ordinary word processors, spreadsheets and case-management software out, while re-capturing them once they are embedded with or functionally dependent upon AI.

Three drafting weaknesses nonetheless merit attention. First, the definition is circular: it defines artificial intelligence partly by reference to software being "functionally dependent upon artificial intelligence," leaving the boundary without an independent anchor at precisely the margin where disputes arise. Second, the operative verbs are conjunctive — "infers, learns, and generates" — which, read literally, would require a system to do all three, though many narrow tools only do one; the OECD text avoids this by using "such as." Third, the listed outputs are "decisions, predictions, and recommendations," omitting content; generative systems that draft text, summaries or pleadings produce content, and a literal reading risks leaving the most common form of judicial generative AI outside the core verb list.

B.4 The DGPSI definition — assessment

DGPSI ties the definition directly to the thing the law actually cares about — the point at which a human ceases to be meaningfully in control. This coheres tightly with the human-primacy principle in Section 4 of the draft and with DGPSI-AI's second principle of one accountable human behind every algorithm. Where the Court reaches the same result through separate provisions on autonomy, risk and the Regulation 20 prohibitions, DGPSI builds the accountability concern into the definition itself.

Its difficulties are of a different character. “Acceptable threshold” is left undefined and is inherently relative, presupposing a standard-setter. The three classes are meant to supply that content, but there is a slight mismatch: Classes 1 and 2 sit on the control axis (self-correction; automatic implementation), whereas Class 3 (creative/affective output) is a capability criterion that need not involve any reduction in human intervention. The classes are also not stated to be hierarchical or mutually exclusive — an agentic generative system could fall in all three — and Class 1’s reference to “code-correcting ability” invites a literalism trap, since most machine learning adjusts weights and parameters rather than rewriting its own code. The intended reading should be made explicit.

B.5 Mapping the two together

The frameworks are complementary rather than contradictory, and the DGPSI classes nest reasonably within the Court’s capability description, while also exposing where that description needs reinforcement:

DGPSI class	Maps to SC capability	Suggested risk tier	Note
Class 1 — self-correcting code without a human developer	Instance of the Court’s “learns.”	Medium–High, depending on the decision it feeds.	Clarify whether parameter/weight adjustment (not only source-code change) is included.
Class 2 — automatically implements a decision affecting a human	“generates decisions” exercised with autonomy.	Critical — and largely the prohibited zone in courts.	Maps onto the Regulation 20 prohibitions (no algorithmic adjudication / automated outcomes).
Class 3 — affective and creative / generative output	Under-specified — “content” is absent from the Court’s verb list.	Medium–High; Critical if outputs enter the record unverified.	Fills the generative-content gap; tie to BSA s.63 disclosure and explainability.

B.6 Recommendation

For an approval gate, the Court’s descriptive definition is good at setting the outer boundary but weak on administrability; DGPSI’s class model is the opposite. The two are therefore best combined:

1. **Retain a descriptive, OECD-aligned definition as the gate** — but cure the three defects: replace the conjunctive “infers, learns, and generates” with “infers, learns, or generates … such as,” add “content” to the listed outputs, and remove the circular reference to “artificial intelligence” in the carve-out (substituting the functional test that the software leverages autonomous learning or probabilistic models to adapt its behaviour).
2. **Layer a control-based classification within the gate** — adopt the Low / Medium / High / Critical risk tiers (Annexure A, Reg. 12), informed by the DGPSI classes, to set the intensity of obligations once a system is within scope.

3. **Require a vendor-confirmed functional test at the approval stage** (Annexure A, Reg. 3(1)(m); DGPSI Deployer MIS-1), and add a definition of "AI Agent" / "agentic AI," presently referred to in Regulation 23(a) but left undefined.

In short: the Court's definition is the better boundary definition — globally aligned and descriptive — while DGPSI's is the better regulatory-trigger definition, fastening onto the loss of human control. Neither is complete alone. Used together, the descriptive definition supplies the category and the threshold-and-class model supplies the graduated response — the architecture a judicial AI regime needs.

B.7 A corrected DGPSI definition

To cure the ambiguities identified above — the undefined "acceptable threshold," the mismatch between the control-based headline and the capability-based classes, the unstated relationship between the classes, and the "code-correcting" literalism — while preserving DGPSI's distinctive accountability-first character, FDPPI proposes the following revised definition.

Core definition (the gate). *An AI System is an automated data-processing system that, for a given input, produces decisions, predictions, recommendations or content which are not fully pre-determined by explicit human-authored instructions, because the system derives or adapts its own processing logic from data, models or probabilistic methods.*

Accountability threshold (when the Standard applies). *A system within the core definition is governed by this Standard where the degree of meaningful human intervention in either (a) the formation of its output, or (b) the application of that output to a decision affecting a person or a business or legal outcome, falls below the accountability threshold. The deployer shall define and record the accountability threshold for each system in its risk documentation. In the absence of such a record, or where the system exhibits any characteristic in Classes 1 to 3 below, the threshold is presumed to be crossed and the system is treated as an AI System requiring governance.*

Classes. *The three classes are independent risk vectors, not an ascending scale, and are not mutually exclusive; a system may fall within more than one, in which case the obligations attaching to each apply cumulatively. Governance intensity (Low / Medium / High / Critical) is fixed by the risk tier, not by the class number.*

- ***Class 1 — Adaptive (self-learning) systems:*** *a system that alters its own decision behaviour — by adjusting parameters, weights, rules, embeddings or operative prompts — without a human developer revising the underlying logic for each such change.*
- ***Class 2 — Autonomous-action (automated-decision) systems:*** *a system whose output is implemented, or applied to a decision affecting a person or a business or legal outcome, without a human being able to review and override that specific output before it takes effect.*
- ***Class 3 — Generative and affective systems:*** *a system that generates novel content (including text, images, audio, video or code), or that infers, simulates or responds to human emotional or behavioural states.*

Interpretation. *A system falls within this Standard on either of two independent grounds: because human control has dropped below the threshold (the control ground — Classes 1 and*

2), or because the system possesses capabilities that generate unknown or emergent risk irrespective of human control (the capability ground — Class 3). Either ground is sufficient on its own.

Short form (for the body of the Standard). *An AI System is an automated data-processing system whose decisions, predictions, recommendations or content are not fully pre-determined by explicit human-authored instructions. It is governed by this Standard where meaningful human intervention in the formation or application of its output falls below a deployer-documented accountability threshold — which is presumed crossed where the system (1) adapts its own decision behaviour without per-change human revision, (2) applies a decision affecting a person without a human able to override the specific output, or (3) generates novel content or infers or responds to human emotional or behavioural states.*

How the revision corrects each ambiguity:

- **Undefined threshold:** made procedurally determinate — the deployer must define and document it per system — backed by a default presumption where it is undocumented or any class characteristic is present, dovetailing with DGPSI's Deviation Justification Document discipline.
- **Headline-versus-classes mismatch:** the definition now rests openly on two grounds — Classes 1 and 2 carry the control axis, Class 3 is placed on a separate capability axis whose application does not depend on any reduction in human intervention.
- **Relationship between classes:** stated to be independent, overlapping and cumulative; the class number denotes risk type, not severity, which is read off the separate risk tier.
- **"Code-correcting" literalism:** Class 1 now refers to adjusting parameters, weights, rules, embeddings or prompts — expressly not limited to rewriting source code — so conventional machine learning, which changes weights rather than code, is plainly captured.

A useful by-product is interoperability: the core-definition sentence mirrors the OECD / EU AI Act / draft Supreme Court descriptive boundary, so the revised DGPSI definition realises the "boundary definition plus control-based classification" synthesis recommended in B.6 above.