

Research Article

Reducing Pendency, Preventing Misuse: AI in India's Judicial System

Arpita Pathak¹

1. Independent researcher

Introduction. India's courts face persistent pendency and misuse of procedural levers. I present a governance-first, judgment-neutral registry-intake assistant that operates on filing metadata only, without predicting case outcomes or recommending sanctions.

Methods. Intake is framed as administrative triage on structured signals (e.g., filing subtype, party/advocate counts, completeness flags, duplicate-hint hashes, and time-to-first-action). A lightweight attention mechanism (logistic regression) surfaces at-risk filings for clerk review; no end-to-end adjudication modeling is used. Evaluation uses a synthetic, CPU-only 12-week workload with weekly clustering and 95% confidence intervals (CIs).

Results. The assistant improved front-end handling rate and reduced processing time at intake; early-adjudgment risk was flagged sooner for clerk attention. Evidence accuracy (flags-only) reached 0.89 (95% CI: 0.88–0.91). Ablations showed incremental gains from (i) registry-rules intake logic, (ii) administrative signals only, and (iii) learned attention; the learned piece remained bounded to metadata.

Discussion. The design aligns with EU AI Act operator duties (risk management, data governance, logging, human oversight) and India's e-Courts Phase III rails, with Digital Personal Data Protection Act, 2023 (DPDP) compliance (data minimization, role-based access, retention). Post-market monitoring and incident reporting are built into the workflow.

Conclusion. A judgment-neutral, metadata-only intake assistant can improve administrative efficiency while avoiding adjudication predictions and automated sanctions, offering a practical approach to responsible AI in justice.

Corresponding author: Arpita Pathak, researcher.arpita@gmail.com

1. Introduction

India's National Judicial Data Grid (NJDG) publishes public dashboards of pendency, disposal, and related indicators, offering an aggregate view of backlogs and operational pressures. These data highlight registry "front-door" frictions—such as incomplete filings, misclassification, and repetitive procedural behaviors—that often propagate avoidable delays into the first listing cycle^{[1][2]}. The approach taken here is deliberately administrative: structure intake, require recorded reasons for system-assisted prompts, surface repeat administrative patterns for oversight, and log each interaction to preserve auditability—while keeping adjudication outside the system's remit. This framing accords with the e-Courts Phase III vision for inclusive, open, and user-centric digitization and with the Digital Personal Data Protection Act, 2023 (DPDP) principles of purpose limitation, data minimization, role-based access, and retention. Given the sensitivity of court systems, the tool is treated as high-risk and mapped to operator duties under the EU Artificial Intelligence Act (AI Act) to ensure risk management, data governance, logging/technical documentation, human oversight, and post-market monitoring are explicit design features rather than afterthoughts^{[3][4][5][6]}.

The work is positioned as judgment-neutral by construction. It does not predict outcomes, rank litigants, or automate sanctions. Its contribution is limited to administrative efficiency and governance: (i) reasons-recorded intake logs (RRIL) to create a per-filing "reasons recorded" trail; (ii) explainability sheets (XS) to communicate what was checked and why; (iii) MPR to register administrative pattern surfaces (e.g., serial early-adjudgment loops); and (iv) post-market monitoring plan (PMMP) to monitor metrics, drill rollback, and document threshold changes. These artefacts operationalize the AI Act's controls, connect to e-Courts Phase III rails, and respect DPDP obligations and constitutional privacy as recognized in *Puttaswamy*^[7].

This paper makes a distinctive contribution by presenting the first governance-first, judgment-neutral AI framework for court intake in India. Unlike prior digitization or e-Courts initiatives that emphasize infrastructure or dashboarding, this design introduces a metadata-only, explainable oversight system aligned with both the EU AI Act's operator duties and India's DPDP 2023 requirements. The novelty lies not in complex machine learning, but in the operational governance artefacts—Reasons Recorded Intake Logs (RRIL), Explainability Sheets (XS), Misuse Pattern Registers (MPR), and Post-Market Monitoring Plans (PMMP) that embed accountability, reversibility, and proportionality directly into registry

workflows. By doing so, the work bridges AI governance theory with deployable practice in a high-risk justice-system context, offering an original, verifiable blueprint for responsible AI adoption.

2. Literature Review

Prior work on court digitization and administrative AI spans digitized filings, queue/triage rules, forms validation, and operational dashboards. In India, NJDG and the e-Courts Phase III vision highlight pendency pressures and registry “front-door” frictions^{[1][4][5]}. Internationally, governance frameworks such as the EU AI Act and NIST AI risk management framework (RMF) emphasize minimal data use, human oversight, logging/technical documentation, and post-market monitoring; explainability guidance (ICO & The Alan Turing Institute) further stresses plain-language reasons for system-assisted actions^{[3][8][9]}. However, a judgment-neutral, metadata-only intake assistant that couples reasons-recorded logs (RRIL), plain-language explainability sheets (XS), a misuse-pattern register (MPR), and a post-market monitoring plan (PMMP)-tailored to Indian rails and DPDP 2023-is largely undocumented. This study addresses that gap by evaluating such a design on synthetic data and mapping operator duties explicitly to deployment artefacts.

The EU AI Act frames a risk-based governance program. Systems used in the administration of justice and democratic processes are treated with heightened scrutiny; regardless of exact categorizations in domestic contexts, a registry-facing intake assistant warrants high-risk operator practices: a maintained risk management system (Art. 9), data governance and quality controls (Art. 10), logging and technical documentation, clear human oversight, and post-market monitoring^[3]. Implementing these duties through operator-side artefacts (RRIL, XS, MPR, PMMP) provides verifiable evidence of compliance and enables consistent audits over time.

India’s e-Courts Phase III program articulates implementation rails for digital filings, registries, and user-centric services. The vision emphasizes inclusivity and openness, enabling integrations that improve administrative throughput without altering adjudication^{[4][5]}. The DPDP Act, 2023 supplies a statutory baseline for lawful processing: notice, purpose limitation, data minimization, role-based access, and time-bound retention^[6]. This legal context informs feature selection (metadata-only) and the avoidance of content-level extraction. In parallel, constitutional privacy doctrine in Justice K. S. Puttaswamy (Retd.) v. Union of India supports the same minimal-data posture and the need for accountability trails^[7]. Where practical guidance is needed, the NIST AI Risk Management Framework 1.0

offers a process model for risk identification, measurement, monitoring, and governance that maps cleanly to the AI Act’s duties^[8].

2.1. Theoretical framework

Registry operations face two recurrent frictions. First, front-door friction: incomplete filings, missing annexures, and misclassification that impede early progress. Second, repeat procedural patterns: serial early-adjudgment loops, duplicate filings, and jurisdictional hopping that consume registry and bench time. NJDG dashboards make these pressures visible in aggregate but do not provide per-filing levers for explanation or audit^{[1][2]}.

The system is constrained to administrative support and governance:

Throughput and quality: increase FEHR (percentage of filings handled at the registry in the first listing cycle) and OT (on-time rectification of returns).

Early risk surfacing: increase EA (agreement between surfaced flags and ground truth) while keeping the cadence of MD (policy-relevant pattern surfaces) conservative.

Auditability: ensure every system-assisted prompt is reasoned (RRIL) and explainable in plain language (XS).

Legality and proportionality: adhere to DPDP minimality, notice, role-based access, and retention; keep adjudication out of scope.

Governance and reversibility: maintain a live risk register and PMMP, version configuration changes, and exercise rollback drills consistent with the AI Act and the NIST AI RMF^{[3][8]}.

3. Research Methodology

3.1. Research design

A corpus of 2,400 filings was generated across 12 synthetic “weeks” (200 per week), with fixed seed 1729.

Variables and priors:

```
case_type ∈ {civil, criminal, writ, misc} with base mix 0.45/0.35/0.15/0.05;  
complex ~ Bernoulli(0.30);  
annexure_missing_count ~ Poisson(0.6);
```

`duplicate_hint ~ Bernoulli(0.08);`

`days_since_prev_listing ~ Log-Normal with $\mu = \ln(14)$, $\sigma = 0.45$.`

A filing counts as front-end-handled if registry checks pass and no avoidable judicial touch occurs in the first listing cycle. A sensitivity run increases complex to 0.50 while holding other distributions fixed. The distributions reflect coarse, nationwide patterns visible on NJDG dashboards and public documentation; parameterization remains conservative to avoid overfitting to any one jurisdiction's mix^[2] (National Judicial Data Grid). (NJDG aggregates are used only for shape; no real records were processed.)^{[1][2]}.

In synthetic evaluation, EA is defined as the agreement between a surfaced flag and the generator's labeled condition (e.g., "mandatory annexure missing" or "duplicate filing hint present"). In deployment, EA would be computed against oversight determinations sampled from RRIL/XS entries.

To reduce concept-piece risk, the generator's collapsed civil/criminal marginals are compared with NJDG national aggregates using χ^2 or Kolmogorov-Smirnov (KS) tests. If divergence is material, proportional reweighting is specified and recorded in the Data Card (kept with the technical file/logs). This procedure uses publicly visible NJDG distributions and program descriptions offered by the Department of Justice; it does not rely on individual-level data^{[1][2]}.

Unit = filing; cluster = week. For FEHR, OT, and EA, I report weekly means with 95% cluster-robust (HC1) confidence intervals using $df = 11$. APC-early is a cluster mean of counts. A 1,000-replicate cluster bootstrap cross-checks HC1; intervals agreed to two decimals. Rounding policy: means to 2 decimals; percentage-point deltas (pp) to 1 decimal to avoid spurious precision. Governance-linked definitions and thresholds are versioned within PMMP; changes require oversight sign-off^{[3][8]}.

3.2. *Participants*

Not applicable. This study uses synthetic, non-identifying data only; no human or animal subjects were involved.

3.2.1. *Instruments*

The author executed all experiments on a laptop-class CPU (8 logical cores, 16 GB RAM; no GPU) using a standard Python stack. A single lightweight model-logistic regression-was used only to surface attention flags from metadata; there was no adjudication prediction or ranking and no external API calls.

The learned-attention variant uses logistic regression (L2; C = 1.0; solver = lbfgs; max_iter = 200) over metadata-only features: `annexure_missing_count`, `duplicate_hint`, `case_type` one-hots, and `days_since_prev_listing`. The XS summarizes salient feature contributions in plain language, following practical guidance on explaining AI-assisted decisions^[9].

Item	Setting
Hardware	Laptop-class CPU (8 logical cores), 16 GB RAM; no GPU
Software	Python toolchain; logistic regression for learned attention (L2; C = 1.0; lbfgs; max_iter = 200)
Data	12 “weeks” × 200 filings/week = 2,400; seed 1729; case_type mix 0.45/0.35/0.15/0.05; complex = 0.30 (sensitivity = 0.50); <code>annexure_missing_count</code> ~ Poisson(0.6); <code>duplicate_hint</code> ~ Bernoulli(0.08); <code>days_since_prev_listing</code> ~ Log-Normal($\mu = \ln(14)$, $\sigma = 0.45$)
Arms	Baseline; Rule-only; RRIL; RRIL + XS; Full (RRIL + XS + MPR + PMMP)
Evaluation	Weekly-cluster means; HC1 95% CIs (df = 11) with cluster-bootstrap cross-check; one-factor sensitivity (complex share); generator-NJDG fit procedure

Table 1. Feature groups and signals used by the registry-intake assistant

Notes: examples include filing subtype, party/advocate counts, completeness flags, duplicate-hint hashes, time-to-first-action, and registry-rule indicators.

3.3. System Overview: Governance-First, Judgment-Neutral

Roles and responsibilities. The Operator (court IT cell or approved integrator) maintains configurations, the risk register, and technical documentation; the Users (registry staff/clerks) review prompts and must approve/decline any suggested action; Oversight (an administrative judge/committee) samples RRIL/XS monthly, signs off threshold changes, and records decisions in quarterly memos. This allocation of duties ensures human-in-the-loop control and creates institutional memory suitable for post-market monitoring under the AI Act^[3].

Core artefacts. Reasons-Recorded Intake Log (RRIL): a per-case log that lists the features viewed, a plain-language reason for any system-assisted prompt, the human approver’s identity, timestamps, and a

reversal link to undo or annotate earlier actions. RRIL is the canonical evidence for logging and oversight audits.

- **Explainability Sheet (XS):** a one-page, plain-language summary of what was checked and why, with a short description of how the human decision was reached and how to contest it. The structure follows practical guidance on explaining AI-assisted decisions to affected persons^[9].
- **Misuse Pattern Register (MPR):** a registry-level list of administrative patterns-e.g., “>2 early adjournments within 30 days”-that warrant policy attention. MPR does not trigger sanctions; it prompts oversight review.
- **Post-Market Monitoring Plan (PMMP):** a governance loop covering metric reviews (FEHR, OT, EA, MD), incident logging, rollback drills, and quarterly oversight memos. PMMP operationalizes Art. 9-10 duties and the Act’s broader post-market posture^[3].
- **Bright-line exclusions.** The tool does not predict outcomes, does not rank litigants or advocates, and does not automate sanctions. These exclusions are necessary to preserve judicial independence and proportionality in processing personal data, consistent with constitutional privacy and the AI Act’s emphasis on human oversight^{[7][3]}.
- **Integration context.** The intake assistant interfaces only with fields already present on e-Courts Phase III rails-filing type/date, party count, representation flag, and annexure checklist-avoiding content extraction. This enables administrative gains while respecting DPDP constraints and reducing integration risk^{[4][5][6]}.

3.4. Metrics and definitions

- **FEHR (Front-End Handling Rate):** proportion of filings fully handled at registry during the first listing cycle.
- **APC-early (Adjournments per Case, first 30 days):** mean adjournments per case within 30 days of initial listing.
- **OT (On-time Rectification):** proportion of returns corrected within the standard rectification window.
- **EA (Escalation Accuracy):** proportion of surfaced flags confirmed against ground truth (synthetic labels here; oversight in deployment).
- **MD (Administrative Misuse Detections):** count of policy-relevant pattern surfaces per week (no sanctions; prompts oversight review).

- These definitions are administrative by construction (judgment-neutral) and map to AI Act logging/oversight expectations^[3] and NIST AI RMF post-market monitoring practice^[8].

3.5. External-validity and uncertainty controls

Generator-NJDG fit. Before reporting results, the generator’s marginals are compared to NJDG aggregates; material divergences trigger proportional reweighting documented in the Data Card. This maintains fidelity to observed filing mixes without ingesting personal data^{[1][2]}.

Uncertainty discipline. I report cluster-robust intervals at the week level, then confirm with a cluster bootstrap. Rounding avoids overstating precision. Definitions and thresholds remain versioned under PMMP, with reversal drills recorded quarterly^{[3][8]}.

Explainability discipline. Each learned-attention prompt generates an XS entry that explains “what was checked and why,” following ICO/Turing guidance to support comprehension and contestation^[9].

3.6. Threats to validity

- **Measurement/reactivity bias.** Awareness of observation can alter clerk behavior. Mitigation: unannounced RRIL/XS sampling; cluster-level estimation.
- **Over-flagging / throughput drag.** Aggressive thresholds can slow the desk. Mitigation: PMMP threshold-pruning gates; monthly oversight review.
- **Filing-mix drift.** Case-type composition may shift. Mitigation: quarterly generator re-fit to NJDG aggregates with documented reweighting in the Data Card.
- **Scope creep into adjudication.** Pressure to expand features to outcomes. Mitigation: bright-line exclusions; governance sign-off; legal basis restricted to administrative metadata (DPDP 2023), with constitutional privacy constraints recognized in jurisprudence.^{[3][6][7][8]}

3.7. Results (synthetic evaluation; 12 weeks; N = 2,400)

3.7.1. Setup summary

All runs were executed on a laptop-class CPU; the learned-attention variant used a single logistic regression over metadata-only features. Arms comprised a status-quo Baseline, a Rule-only baseline (annexure/duplicate rules without learned attention), Ablation-A (RRIL), Ablation-B (RRIL + XS), and Full (RRIL + XS + MPR + PMMP). Evaluation used weekly clusters (N = 12) with HC1 95% confidence intervals;

a 1,000-replicate cluster bootstrap cross-check agreed to two decimals. The generator’s proportions were designed to reflect coarse, nationwide patterns visible on NJDG aggregates; sensitivity increased the complex-filing rate from 0.30 to 0.50. See Table 1.

3.7.2. Primary outcomes

Relative to Baseline, Full improved registry-level administrative measures:

- FEHR increased 0.63 → 0.79 ($\Delta = +16.3$ pp; 95% CI for Full: 0.78-0.80).
- APC-early decreased 1.80 → 1.40 ($\Delta = -0.40$; 95% CI for Full: 1.39-1.40).
- OT increased 0.53 → 0.74 ($\Delta = +21.6$ pp; 95% CI for Full: 0.72-0.76).
- EA (flags only) reached 0.89 (95% CI: 0.88-0.91).
- MD stabilized at ≈ 1 policy-relevant pattern surface per week under Full (conservative surfacing; mandatory human approval).
- Ablations showed stepwise gains: RRIL improved FEHR/OT relative to Baseline; adding XS lifted EA and further improved FEHR/OT. The Rule-only baseline delivered modest improvements over Baseline; learned attention added a further step-up preserved through RRIL, RRIL + XS, and Full.

Metric	Baseline	Rule-only	RRIL	RRIL + XS	Full
FEHR	0.63 [0.60-0.65]	0.66 [0.64-0.68]	0.69 [0.67-0.71]	0.75 [0.73-0.76]	0.79 [0.78-0.80]
APC-early	1.80 [1.79-1.80]	1.76 [1.75-1.76]	1.75 [1.75-1.75]	1.60 [1.59-1.60]	1.40 [1.39-1.40]
OT	0.53 [0.51-0.55]	0.57 [0.55-0.59]	0.61 [0.59-0.64]	0.71 [0.68-0.73]	0.74 [0.72-0.76]
EA	-	-	0.82 [0.80-0.84]	0.86 [0.84-0.88]	0.89 [0.88-0.91]
MD (per week)	0.00 [-]	0.20 [0.00-0.40]	0.50 [0.17-0.83]	0.92 [0.73-1.10]	1.00 [1.00-1.00]

Table 2. Results (weekly-cluster means; 95% CIs)

Notes: FEHR = Front-End Handling Rate; APC-early = Adjournments per Case in first 30 days; OT = On-time Rectification; EA = Escalation Accuracy; MD = Administrative Misuse Detections per week.

3.7.3. Sensitivity to complexity mix

When the share of “complex” filings increased from 0.30 to 0.50, Full experienced modest degradations: FEHR -1.7 pp, OT -4.7 pp, APC-early +0.05, EA -0.5 pp. Effects remained within monitoring tolerances and did not invert the ablation ranking.

3.7.4. Generator-NJDG fit check (procedural)

Collapsed civil/criminal marginals are compared against NJDG aggregates using χ^2 or KS tests; any material divergence is addressed by proportional reweighting recorded in the Data Card. No individual-level data are used; the check enforces alignment with observed mixes^{[1][2]}.

3.8. Discussion

3.8.1. Administrative levers, not adjudication

Gains derive from **administrative discipline** rather than predictive adjudication. RRIL elevates throughput by making overrides reason-accountable; XS increases EA by aligning surfaced features with human expectations, reducing unnecessary escalations. The MPR and PMMP in Full introduce a conservative surfacing cadence (MD \approx 1/week) and a continuous improvement loop without automated sanctions. This posture aligns with EU AI Act operator duties—documented risk management, data governance, logging, human oversight, and post-market monitoring—while remaining compatible with e-Courts Phase III operational rails and DPDP minimality^{[3][4][5][6]}.

3.8.2. Role of explainability and logging

The obligation to **explain** what was checked and why—operationalized via XS—addresses transparency requirements and supports contestation, consistent with practical guidance^[9]. RRIL provides the audit trail regulators expect: a human approver, a timestamp, and a reversibility path. Together they satisfy the **logging/technical-file** expectations under the Act and help institutions demonstrate proportionality and accountability^[3].

3.8.3. Learned attention versus rules

The **Rule-only** baseline improves FEHR and OT relative to status-quo practices, indicating checklists alone have value. The learned-attention layer contributes additional-but bounded-gains, supporting the

choice of a simple, interpretable model over opaque predictors. This aligns with the **judgment-neutral** stance and the **risk-minimization** required for justice-system contexts^[3].

3.8.4. Sensitivity and operational resilience

The single-factor sensitivity demonstrates expected degradations when filings are more complex. The EA decline is minor, suggesting that the explanation/oversight loop is resilient. Under a PMMP, this would prompt threshold adjustments and targeted staff guidance rather than architectural changes^{[3][8]}.

3.8.5. Threats to validity and mitigations

Potential reactivity and over-flagging are contained by cluster-level estimation and PMMP threshold-pruning gates. Filing-mix drift is addressed by periodic generator reweighting against NJDG aggregates with changes versioned in the Data Card. Scope creep is constrained by bright-line exclusions and oversight sign-off^{[3][1][8]}.

3.8.6. Policy and deployment implications

Because the approach uses metadata-only features and a human-approval gate for any action, deployment can proceed under DPDP minimality and e-Courts Phase III rails without scraping or third-party enrichment. The result is a tractable institutional pathway: measurable administrative improvements, a verifiable compliance posture, and a reversible integration design^{[4][5][6]}.

Generalizability. Although evaluated on a registry-intake assistant, the governance pattern-judgment-neutral scope, metadata-only processing, RRIL/XS logging, and PMMP oversight-transfers to other public-sector AI that supports administrative decisions (licensing, benefits intake, case triage). Transfer is appropriate when (i) tasks are administrative (not adjudicative), (ii) lawful metadata suffices, (iii) human approval is required, and (iv) post-market monitoring and rollback are feasible.

3.9. Implementation Blueprint for Courts

Phase 0 - Readiness (Week 0).

- Approve the scope document (judgment-neutral; metadata-only).
- Create the **risk register** and **Data Card**; appoint an oversight committee.

Phase 1 - Baseline and templates (Weeks 1-2).

- Baseline FEHR, OT, APC-early using existing standard operating procedures (SOPs).
- Finalize **RRIL** and **XS** templates; publish notices; configure access controls and retention schedules consistent with **DPDP**.

Phase 2 - RRIL live (Ablation-A) (Weeks 3-6).

- Enforce reasons-recorded logging on all triage suggestions.
- Monitor FEHR and OT weekly; convene oversight for log sampling.

Phase 3 - RRIL + XS (Ablation-B) (Weeks 7-8).

- Attach **XS** to each system-assisted suggestion.
- Track **EA** against synthetic ground truth in the sandbox; prepare oversight sampling criteria for deployment.

Phase 4 - Full governance (Weeks 9-12).

- Enable **MPR** and **PMMP**.
- Conduct one rollback drill; document any threshold proposals and approvals; record a brief quarterly memo.

3.9.1. Data protection and legality

- Limit processing to registry-lawful metadata (purpose limitation, data minimization, role-based access, time-boxed retention) under **DPDP**.
- Keep adjudication out of scope; no ranking; no automated sanctions^{[3][6]}.

3.9.2. Success criteria

- **Primary:** FEHR and OT increase relative to Baseline; APC-early decreases; $EA \geq 0.85$ with conservative surfacing.
- **Secondary:** MD \approx 1/week with demonstrated policy relevance; zero unresolved incidents; successful rollback drill.

3.9.3. Operational documentation

RRIL/XS snapshots archived; Data Card and risk register updated monthly; PMMP incident log and threshold-change proposals versioned and signed by oversight^{[3][8]}.

3.10. Compliance Mapping to the EU Artificial Intelligence Act

Regulatory posture. The intake assistant is operated as a **high-risk** system due to its proximity to justice administration. Operator duties are implemented through artifacts and procedures described earlier (RRIL, XS, MPR, PMMP).

Art. 9 - Risk management. A maintained **risk register** documents hazards, mitigations, residual risks, owners, and review dates across releases. **Rollback drills** and incident logging are scheduled under the PMMP. Risks specific to over-flagging, filing-mix drift, and scope creep are tracked with triggers for threshold pruning and oversight sign-off^[2].

Art. 10 - Data and data governance. A **Data Card** records synthetic-data provenance, sampling and reweighting against NJDG aggregates (External Validity procedure), bias probes, refresh cadence, and validation notes. No personal data are ingested for the experiments; deployment restricts processing to registry-lawful **metadata** (filing type/date, party count, representative flag, annexure checklist)^[2].

Logging & technical documentation. RRIL is the canonical per-case log; a design dossier versions configurations, thresholds, and changes. XS provides plain-language explanations attached to case files to support audit and contestation.

Human oversight. All prompts require **approve/decline** by registry staff; Oversight (administrative judge/committee) samples RRIL/XS monthly and approves threshold changes.

Post-market monitoring. The PMMP defines metric review (FEHR, OT, APC-early, EA, MD), incident handling, threshold-change governance, and rollback practice with quarterly oversight memos^[2].

3.11. Ethics, Privacy, and Legal Commitments

Judgment-neutral scope. The tool **does not** predict adjudicatory outcomes, rank litigants, or impose automated sanctions, preserving judicial independence and proportionality.

DPDP compliance. Processing rests on a lawful basis available to court administration; **purpose limitation, data minimization, role-based access, and time-bound retention** are enforced. The integration relies on **e-Courts Phase III** fields only; no content extraction or third-party enrichment is used^{[6][4][5]}.

Constitutional privacy. The design aligns with the privacy guarantees articulated in *Justice K. S. Puttaswamy (Retd.) v. Union of India*, emphasizing minimality, accountability, and auditability^[7].

Practical governance scaffolding. The NIST AI RMF 1.0 is used for process discipline (risk identification, measurement, governance), complementary to the EU AI Act. **ICO/Turing** guidance informs the content and tone of **XS** so explanations are understandable to affected persons^{[8][9]}.

Use of AI tools. A language model (ChatGPT) was used only for language editing. The author reviewed and is fully responsible for all content.

4. Conclusion

A governance-first, judgment-neutral intake system that structures filings, explains prompts, logs actions, and monitors effects yields consistent gains on administrative metrics in a CPU-only synthetic evaluation, while conforming to operator duties under the EU AI Act and aligning with e-Courts Phase III rails and DPDP 2023. The appropriate next step is a controlled pilot on real workflows under the same guardrails, with oversight sampling and PMMP drills maintained throughout^{[3][4][5][6]}.

Recommendations

These recommendations apply to the intake assistant and to comparable administrative AI tools that meet the same guardrails (administrative scope, metadata-only processing, human approval, RRIL/XS logging, PMMP oversight)

- **R1.** Maintain judgment-neutral scope: administrative intake only; **no** adjudication prediction/ranking/sanctions. Any scope change requires oversight approval and a recorded data protection impact assessment (DPIA) addendum.
- **R2.** Limit processing to registry-lawful **metadata only** (filing type/date, party count, representation flag, annexure checklist). No content extraction; no third-party enrichment.
- **R3.** Proceed to scale-up only if, versus Baseline: **FEHR ≥ 10 pp, OT ≥ 15 pp, EA ≥ 0.85 , APC-early ≥ 0.30** . All deltas computed as weekly cluster means with 95% CIs.
- **R4.** Monthly RRIL/XS sampling; **quarterly** oversight memo summarizing metrics, incidents, and threshold changes. Changes must cite evidence and be versioned in the design dossier.
- **R5.** Define a ≤ 1 **business day** rollback service-level agreement (SLA) for any adverse effect; test via a quarterly drill and archive drill reports in the PMMP.
- **R6.** Compare filing-mix marginals with NJDG each quarter; if divergence is material, apply proportional re-weighting and record in the Data Card.

- **R7.** Publish a short notice on what is checked, why, and how to seek human review/correction; train registry staff on **RRIL** and **XS**.
- **R8.** Approve scale-up only when: targets in **R3** are met for **≥4 consecutive weeks**, zero unresolved incidents, last rollback drill passed, and the latest oversight memo is signed.

Statements and Declarations

Policy Significance

This article presents a judgment-neutral, metadata-only intake assistant that improves registry throughput without predicting outcomes or automating sanctions. The framework operationalizes operator-duty controls from the EU AI Act and India’s DPDP 2023 via Reasons-Recorded Intake Logs (RRIL), Explainability Sheets (XS), a Misuse Pattern Register (MPR), and a Post-Market Monitoring Plan (PMMP), embedding auditability, reversibility, and proportionality into routine workflows. In a 12-week synthetic evaluation, I observe higher front-end handling and on-time rectification, fewer early adjournments, and high escalation accuracy. The design runs on e-Courts Phase III rails without content scraping or third-party enrichment and generalizes to other public-sector intake tasks (e.g., licensing, benefits). It provides a verifiable, deployable path to responsible AI adoption in courts. Although evaluated for India’s courts, the governance pattern is jurisdiction-agnostic and transfers to other public-sector intake workflows (e.g., licensing, benefits) under similar operator-duty regimes (EU AI Act, DPDP-style statutes).

Funding

No funding was received for this work.

Conflicts of Interest

The authors declare that they have no conflicts of interest to this work.

Ethics

This study uses synthetic, non-identifying data and restricts processing to administrative metadata. No human or animal subjects were involved. The system excludes adjudication prediction, ranking, and automated sanctions, and aligns with DPDP 2023 principles and e-Courts Phase III rails. The same

governance posture (minimal data, human oversight, logging, reversibility) should be applied to comparable administrative AI deployments.

Data Availability

This study uses synthetic, nonidentifying data. All replication materials are openly available at DOI: <https://doi.org/10.5281/zenodo.17762727>. The deposit includes the fixed-seed dataset, the synthetic generator and configuration (seed 1729), the analysis script implementing the metadata only logistic-regression attention layer, a README with exact run steps, and the RRIL, XS, MPR, and PMMP templates. Licenses: code MIT, data CCO 1.0, templates CC BY 4.0.

Author Contribution

Arpita Pathak: Conceptualization, Methodology, Software, Data curation, Formal analysis, Investigation, Validation, Resources, Visualization, Writing - original draft, Writing - review & editing, Project administration.

Acknowledgements

The author thanks reviewers and colleagues for constructive feedback on early drafts.

References

1. ^{a, b, c, d, e, f, g, h}National Judicial Data Grid (NJDG) (n.d.). "NJDG Dashboards." National Judicial Data Grid (NJDG). <https://njdg.ecourts.gov.in/>.
2. ^{a, b, c, d, e, f, g}Department of Justice (2025). "The National Judicial Data Grid (NJDG)." Department of Justice. <https://doj.gov.in/the-national-judicial-data-grid-njdg/>.
3. ^{a, b, c, d, e, f, g, h, i, j, k, l, m, n, o, p, q, r, s, t, u, v}European Union (2024). "Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)." Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>.
4. ^{a, b, c, d, e, f, g, h}e-Committee, Supreme Court of India (2022). "Vision Document for Phase III of e-Courts Project." e-Committee, Supreme Court of India. <https://ecommitteesci.gov.in/document/vision-document-for-phase-iii-of-ecourts-project/>.

5. ^{a, b, c, d, e, f, g, h}e-Committee, Supreme Court of India (2023). "Vision Document for Phase III of e-Courts Project (Overview)." e-Committee, Supreme Court of India. <https://ecommitteesci.gov.in/vision-document-for-phase-iii-of-ecourts-project/>.
6. ^{a, b, c, d, e, f, g, h, i}Ministry of Electronics and Information Technology (2023). "The Digital Personal Data Protection Act, 2023 (No. 22 of 2023)." Ministry of Electronics and Information Technology. <https://www.meity.gov.in/static/uploads/2024/06/2bf1f0e9f04e6fb4f8fef35e82c42aa5.pdf>.
7. ^{a, b, c, d, e}Supreme Court of India (2017). "Justice K. S. Puttaswamy (Retd.) v. Union of India." Supreme Court of India. (2017) 10 SCC 1. https://www.scobserver.in/wp-content/uploads/2021/10/1-266Right to Privacy_Puttaswamy_Judgment-Chandrachud.pdf.
8. ^{a, b, c, d, e, f, g, h, i, j, k}National Institute of Standards and Technology (2023). "Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1)." National Institute of Standards and Technology. doi:[10.6028/NIST.AI.100-1](https://doi.org/10.6028/NIST.AI.100-1).
9. ^{a, b, c, d, e, f}Information Commissioner's Office, The Alan Turing Institute (2020). "Explaining Decisions Made with AI." Information Commissioner's Office. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/explaining-decisions-made-with-artificial-intelligence/>.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.