DOI: 10.71317/jgst.2.11.2026.1028
Journal of Global Social Transformation
Journal homepage: https://rjsaonline.org/index.php/JGST

Institutional Readiness before Artificial Intelligence: A Governance Framework for AI-Assisted Judicial Reform in Pakistan

Syeda Masheea Fatima[1]
[1] Change Management & Communication Specialist, Supreme Court of Pakistan; Independent Researcher, Islamabad, Pakistan
Email: masheeafatima@gmail.com

ARTICLE INFO

Submission:
October 08, 2026
Revised:
October 30, 2026
Accepted:
November 04, 2026
Available Online:
November 12, 2026

Keywords:
Artificial Intelligence; Judicial Reform; Institutional Readiness; AI Governance; Public Policy; Pakistan.

Corresponding author:
Syeda Masheea Fatima
Email:
masheeafatima@gmail.com

ABSTRACT

Courts around the world are facing increasing pressure to improve their performance without compromising impartiality or judicial independence. Artificial intelligence is increasingly being touted as part of the solution. Most governance frameworks concentrate on the responsible use of AI, particularly transparency, accountability and human oversight. Pakistan's National Guidelines for the Use of Artificial Intelligence in Judicial Institutions take a similar approach. What remains less developed is the question of institutional readiness: what should already be working within a court before AI is introduced?

Pakistan offers a useful setting in which to examine that question. Judicial reform, digital transformation and AI governance are unfolding at the same time. This overlap makes Pakistan particularly useful for examining whether institutional reform is keeping pace with technological change. This study uses comparative qualitative policy analysis to examine Pakistan's National Guidelines alongside the OECD AI Principles, the UNESCO Recommendation on the Ethics of Artificial Intelligence and selected international frameworks for judicial AI. The comparison focuses on the institutional assumptions behind these documents, particularly what they expect courts to have in place before their principles can work in practice.

The frameworks explain how AI should be used once it has been introduced, yet they offer far less guidance on the institutional arrangements that should already exist before that point. Writing these principles into policy is one thing. Putting them into practice is another. Courts need clear responsibilities, reliable administrative systems and ways of explaining how AI is being used. Without these institutional arrangements, ethical principles risk remaining policy statements instead of changing how the justice system actually works.

Despite differences in language and emphasis, similar institutional foundations appeared across almost all of the documents reviewed. Transparent administration, accountable decision-making, human oversight, reliable digital governance and sustained public engagement appear again and again. Together, these features offer a practical way of judging whether a court is actually ready for AI. Looking only at technical capability tells us very little. Technical safeguards such as explainable AI still matter, but they cannot substitute for institutions that lack credibility or public trust.

The implication for judicial reform is practical. Before introducing AI, courts should examine whether their existing administrative systems can support transparent, accountable and properly supervised use. Institutional readiness should therefore be assessed before technological readiness. A court may have the technical capacity to use AI without yet having the institutional capacity to govern it well.

Introduction

Artificial intelligence is beginning to find a place in court administration. Legal research, document processing, case management and other routine functions are obvious areas of interest, particularly in judicial systems dealing with large caseloads and persistent delays. The attraction is easy to understand. If some administrative work can be done faster, courts may be able to use their limited time and resources more effectively. But courts cannot judge a technology only by whether it saves time. Fairness, judicial independence and the ability to account for how decisions are made remain part of the equation. Reiling makes this point in a different way by showing that the possibilities and risks of judicial AI vary considerably according to the kind of court function involved.[1]

Much of the policy response has therefore concentrated on responsible use. Transparency, accountability, privacy, non-discrimination and human oversight now appear across major AI governance frameworks. UNESCO's *Recommendation on the Ethics of Artificial Intelligence* provides the broader international foundation, while the CEPEJ Ethical Charter and UNESCO's more recent guidelines for courts and tribunals address these concerns more directly in judicial settings.[2] The field has also moved beyond broad ethical statements. CEPEJ's operational assessment tool, for example, allows safeguards to be considered before an AI system is acquired or developed.[3]

Pakistan entered this discussion more formally in 2026 with the adoption of the *National Guidelines for Use of Artificial Intelligence in Judicial Institutions of Pakistan*. The Guidelines cover superior and subordinate courts as well as other judicial institutions and envisage AI in areas such as legal research, case management and administrative work. Human judgment remains central: AI is intended to assist rather than replace judicial decision-making. The Guidelines also address accountability, privacy, security, bias and capacity-building.[4]

This development did not take place in isolation. Pakistan's courts were already moving through a wider process of digital reform, including e-filing and other technology-enabled judicial services. AI therefore enters an institution that is changing in several ways at once. This makes Pakistan a useful case, but it also complicates the question of readiness. A court may have digitised records and access to AI tools without necessarily having settled who is responsible for their use, how outputs will be checked, what information should be disclosed, or what happens when an AI-supported process goes wrong.

These concerns are already visible in Pakistan. Following engagement with judicial academies in 2026, UNESCO reported the use of AI tools for legal research, drafting and case management, alongside uneven institutional preparedness. Training, explainability, bias and clearer institutional guidance emerged as practical concerns.[5] This suggests that the relevant policy problem is changing. Pakistan now has principles for responsible judicial AI. The question is how those principles become workable institutional arrangements.

That distinction provides the starting point for this paper. Technological readiness and institutional readiness are related, but they are not the same thing. A court may have the infrastructure and technical capacity to use an AI application while lacking the administrative arrangements required to govern it properly. Recent work on AI adoption in public institutions has begun to make a similar distinction, shifting attention from the readiness of the technology to the readiness of the organisation expected to receive it.[6] Courts, however, require a more specific treatment because questions of institutional capacity intersect with judicial independence, procedural fairness, accountability and public legitimacy.

This paper therefore asks two questions: what institutional conditions should be assessed before AI is deployed within a judicial institution, and how far do existing Pakistani and international governance frameworks help courts make that assessment? It uses comparative qualitative policy analysis to examine Pakistan's National Guidelines alongside selected UNESCO, OECD and CEPEJ

[1] A. D. (Dory) Reiling, "Courts and Artificial Intelligence," *International Journal for Court Administration* 11, no. 2 (2020): 1–10, https://doi.org/10.36745/ijca.343.

[2] UNESCO, Recommendation on the Ethics of Artificial Intelligence (Paris: UNESCO, 2021); European Commission for the Efficiency of Justice (CEPEJ), European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and Their Environment (Strasbourg: Council of Europe, 2018); Juan David Gutiérrez, Guidelines for the Use of AI Systems in Courts and Tribunals (Paris: UNESCO, 2025).

[3] CEPEJ, Assessment Tool for the Operationalisation of the European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and Their Environment (Strasbourg: Council of Europe, 2023).

[4] National Judicial (Policy Making) Committee, National Guidelines for Use of Artificial Intelligence in Judicial Institutions of Pakistan (Islamabad: Law and Justice Commission of Pakistan, 2026).

[5] UNESCO, "AI and Courts: A Tale of Two Intelligent Entities," July 14, 2026.

[6] Erika Fille Legara, Elmo Domino Jose, and Paula Joy Martinez, "Beyond Model Readiness: Institutional Readiness for AI Deployment in Public Systems," paper presented at the 2nd Workshop on Technical AI Governance Research, ICML 2026 (2026).

instruments. From that comparison, the paper develops an Institutional Readiness Framework for Judicial AI around five areas: administrative transparency, accountability architecture, human oversight capacity, digital and data governance, and public legitimacy and engagement. The purpose is not to produce another set of ethical principles. It is to provide a way of judging whether the institution expected to implement those principles is in a position to do so.

Judicial AI and the Institutional Readiness Gap

From judicial technology to judicial governance

Research on technology in courts initially dealt largely with what digital systems could do and where they could reasonably be introduced. Reiling's work is useful here because it separates different forms of judicial work rather than treating "AI in courts" as a single activity. Routine and highly structured tasks are different from work requiring interpretation, discretion or adjudication. Susskind reaches the issue from a broader perspective, arguing that technological change can alter not only the speed of court processes but the way justice itself is organised and accessed.[7]

As AI has become more capable, the literature has moved towards questions of governance. This is particularly important because inserting a human somewhere in an automated process does not necessarily preserve accountability. Research on automation and the rule of law has shown that technological systems can make it harder to identify how a decision was reached, where responsibility lies and how an affected person can challenge it.[8] Those problems become especially sensitive in courts. Judicial decisions are expected to be reasoned and open to scrutiny, and the use of AI should not make the exercise of judicial authority harder to understand or review.

This concern also has a public-facing dimension. Trust in courts does not depend only on whether a person wins or loses a case. Procedural-justice research has long shown that people also respond to the fairness of the process, whether they believe they have been heard and whether authority appears to have been exercised neutrally.[9] More recent research on algorithms in courts suggests that the stage at which AI enters judicial decision-making matters as well. Barysė and Sarel found different public responses to automation depending on whether an algorithm was involved in gathering information or in later stages of judicial decision-making.[10] This makes the institutional setting of AI important in its own right.

Principles are not the same as capacity

The main governance frameworks respond to many of these concerns. Human oversight, transparency and accountability are now well-established requirements. Yet the existence of a principle says relatively little about an institution's ability to implement it.

Human oversight illustrates the problem. A policy may require a judge or court official to remain responsible for an AI-supported process. For that responsibility to mean anything in practice, the person concerned must know when AI has been used, understand enough about its function to question an output and have the authority to disregard it. The court also needs procedures for recording errors and deciding what happens when a system repeatedly produces unreliable results. Human oversight, in other words, rests on human capacity and administrative arrangements.

Transparency has the same institutional dimension. Much of the AI literature understandably concentrates on explainability: whether the operation or output of an AI system can be made intelligible. Judicial transparency, however, also involves more ordinary questions. Does the court keep a record of where AI is being used? Can it explain the purpose for which a system was introduced? Is responsibility for the system identifiable? What information can a litigant obtain if its use affects a process concerning their case? Scholarship on judicial AI has increasingly connected these questions with the continuing requirements of transparency and accountability in judicial decision-making.[11]

Data governance raises another difficulty. AI systems depend on data, but judicial records are not simply technical inputs. They may contain personal or sensitive information, differ in quality between courts and have been created for purposes that pre-date

[7] Reiling, "Courts and Artificial Intelligence"; Richard Susskind, Online Courts and the Future of Justice (Oxford: Oxford University Press, 2019).
[8] Monika Zalnieriute, Lyria Bennett Moses, and George Williams, "The Rule of Law and Automation of Government Decision-Making," Modern Law Review 82, no. 3 (2019): 425–55, https://doi.org/10.1111/1468-2230.12412.
[9] Tom R. Tyler, "Procedural Justice and the Courts," Court Review 44, nos. 1–2 (2007–8): 26–31.
[10] Dovilė Barysė and Roee Sarel, "Algorithms in the Court: Does It Matter Which Part of the Judicial Decision-Making Is Automated?," Artificial Intelligence and Law 32 (2024): 117–46, https://doi.org/10.1007/s10506-022-09343-6.
[11] Tatiana Dancy and Monika Zalnieriute, "AI and Transparency in Judicial Decision Making," Oxford Journal of Legal Studies 46, no. 1 (2026): 1–34, https://doi.org/10.1093/ojls/gqaf030.

AI. Privacy, security, access controls, retention and the quality of digitised records therefore become part of the institution's ability to use AI responsibly. OECD work on AI in government and justice similarly treats data, governance capacity, impact assessment and continuing oversight as part of responsible public-sector adoption rather than as separate technical questions.[12]

These examples point to a difference between having an AI governance policy and having the institutional capacity to carry it out. Pakistan's National Guidelines illustrate this well. They establish principles of human oversight, accountability, privacy and capacity-building, but those principles must ultimately be implemented by judicial institutions with different administrative structures and levels of technological development.[13] The policy can establish the standard; the institution has to make it work.

What is actually missing?

The research gap needs some care because existing frameworks have already moved towards implementation. It would be inaccurate to argue that they deal only with AI after it has been introduced. CEPEJ's Assessment Tool can be used before acquisition or development, and UNESCO's court-specific guidance covers the adoption, use and oversight of AI across judicial institutions.[14]

Nor is institutional readiness entirely new to AI research. Legara, Jose and Martinez argue that technical viability should be separated from the readiness of the public institution receiving an AI system. Their work identifies organisational conditions that can prevent apparently viable technology from being implemented responsibly.[15] This is an important development because it moves the discussion beyond the properties of the AI system itself.

The gap becomes clearer when this idea is brought into the judicial setting. Existing judicial AI frameworks contain many elements of institutional readiness, but these elements are scattered across requirements concerning ethics, data, human oversight, procurement, risk and training. They generally do not bring those requirements together around a single prior question: what should a court be able to demonstrate about itself before a particular AI use proceeds?

Courts also complicate a general public-sector model of readiness. Accountability does not operate exactly as it does in an ordinary government department. Judges require decisional independence, while court administrations, registries and technical units carry different responsibilities. The legitimacy of judicial authority also depends heavily on procedure and reason-giving. A technically effective AI application can therefore create institutional difficulties even when its performance is satisfactory.

For this paper, institutional readiness for judicial AI means the demonstrated capacity of a judicial institution to govern a particular AI use through functioning arrangements for transparency, accountability, human oversight, data governance and public legitimacy. The word in particular matters. Readiness should not be treated as a general status that a court either possesses or lacks.

A court might reasonably be prepared to use AI for transcription while remaining unprepared to use it for a function that could influence judicial reasoning. The level of oversight, explanation and review required in the two situations would not be the same. Readiness can also change. Staff leave, systems are replaced, new risks emerge and an application may begin to be used differently from the purpose for which it was originally approved.

For that reason, the framework developed here does not classify a judiciary simply as "AI-ready." It asks whether the necessary institutional conditions exist for a specified use, and whether that use should be rejected, piloted under safeguards or allowed to proceed through controlled deployment. This moves institutional readiness closer to an actual governance decision rather than another statement of principle.

Pakistan's Judicial AI Context

Pakistan's move towards judicial AI is part of a wider process of digital reform rather than a standalone technology initiative. In May 2025, the National Judicial Automation Committee (NJAC) was reconstituted with representation from the Supreme Court, Federal Shariat Court, High Courts, the Ministry of Information Technology and Telecommunication, the National Information Technology Board and the Law and Justice Commission of Pakistan. Its mandate placed automation within a broader effort to modernise judicial administration and improve coordination across institutions.[16] This matters for the present study because AI is

[12] OECD, Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions (Paris: OECD Publishing, 2025), https://doi.org/10.1787/795de142-en.

[13] National Judicial (Policy Making) Committee, National Guidelines for Use of Artificial Intelligence in Judicial Institutions of Pakistan.

[14] CEPEJ, Assessment Tool; Gutiérrez, Guidelines for the Use of AI Systems in Courts and Tribunals.

[15] Legara, Jose, and Martinez, "Beyond Model Readiness."

[16] Supreme Court of Pakistan, "National Judicial Automation Committee Resolves to Push Digital Transformation of Judiciary," press release, May 7, 2025.

entering a system in which digital infrastructure, court processes and institutional arrangements are themselves still being developed.

That reform process became more concrete during 2025–26. The National Judicial (Policy Making) Committee approved the introduction of e-filing in district courts as the first phase of a national e-court project and directed the High Courts to review procedural rules required for electronic filing. Importantly, the decision called for a phased approach that took account of system preparedness and capacity-building at district level.[17] Digitalisation is therefore not occurring evenly across a single administrative structure. Pakistan's judicial system comprises institutions with different resources, procedures and levels of technological development. Any discussion of AI readiness has to take this variation seriously.

The *National Guidelines for Use of Artificial Intelligence in Judicial Institutions of Pakistan* emerged against this background. Developed through the NJAC following consultation with the High Courts and relevant experts and institutions, the Guidelines were approved by the NJPMC in 2026. They envisage AI as an assistive technology for functions including legal research, case management and document processing while retaining human judgment as the basis of judicial decision-making. Accountability, transparency, fairness, protection against bias, data security and capacity-building are all addressed.[18] The approach is cautious rather than technology-led: AI is permitted to support judicial work, but not to displace judicial responsibility.

There is also an important institutional feature in the way the Guidelines have been designed. They provide a national baseline but do not prescribe an identical implementation model for every court. High Courts retain the authority to determine arrangements according to their institutional requirements.[19] This reflects the constitutional and administrative structure of Pakistan's judiciary, but it also creates a practical governance question. If implementation is adapted at institutional level, readiness cannot simply be inferred from the existence of national guidelines. Each institution must have the capacity to translate those principles into working procedures.

Developments on the ground suggest that this issue is already relevant. UNESCO's consultations with judicial academies across Pakistan in June 2026 found that judges and judicial officers were using AI tools for legal research, drafting and case management. The consultations also identified uneven preparedness. Participants raised the need for training, clearer guidance and better understanding of algorithmic bias and explainability. Of particular relevance to this study, judicial academies emphasised that existing knowledge and competencies should first be assessed before capacity-building interventions were designed.[20] The finding points to a broader issue: adoption can begin informally, through the choices of individual users, before an institution has fully developed the arrangements needed to govern that use.

Pakistan therefore presents an unusual policy moment. It would be inaccurate to describe the judiciary as entering AI without governance. National principles now exist, judicial institutions are discussing capacity-building, and digital reforms such as e-filing are progressing. At the same time, these developments do not establish that every court, or every proposed AI application, is institutionally ready. The more useful question is how the principles already adopted can be translated into evidence of preparedness at institutional level.

This distinction also prevents the paper from turning into a criticism of Pakistan's National Guidelines. The Guidelines perform a different function. They establish the principles within which judicial AI should develop. The framework proposed in this paper takes the next step by asking how a court can determine whether it currently possesses the administrative, human and governance capacity required to put those principles into practice. In this sense, institutional readiness is better understood as an operational complement to Pakistan's emerging judicial AI framework rather than an alternative to it.

Research Design and Method

This study uses comparative qualitative policy analysis to examine how institutional readiness is addressed within existing frameworks for judicial and public-sector AI governance. Pakistan serves as the primary policy case, while international instruments are used comparatively to identify institutional requirements that recur across different approaches to responsible AI. The purpose

[17] National Judicial (Policy Making) Committee, decision on "Digital Transformation Roadmap—Launch of E-Filing of Cases," 2026. The Committee approved e-filing across district courts and called for a phased approach addressing procedural clarity, system preparedness and capacity-building.

[18] National Judicial (Policy Making) Committee, National Guidelines for Use of Artificial Intelligence in Judicial Institutions of Pakistan (Islamabad: Law and Justice Commission of Pakistan, 2026).

[19] National Judicial (Policy Making) Committee, National Guidelines. The national framework preserves the administrative and judicial autonomy of the High Courts in determining appropriate implementation arrangements.

[20] UNESCO, "AI and Courts: A Tale of Two Intelligent Entities," July 14, 2026.

is not to rank the frameworks or measure their legal strength. Instead, the comparison asks what each framework expects judicial institutions to be capable of doing before, and during, the introduction of AI.

Four bodies of policy material form the core of the analysis: Pakistan's *National Guidelines for Use of Artificial Intelligence in Judicial Institutions*; UNESCO's *Recommendation on the Ethics of Artificial Intelligence* and its subsequent *Guidelines for the Use of AI Systems in Courts and Tribunals*; the OECD AI Principles and related work on AI in justice administration; and the CEPEJ Ethical Charter together with its operational assessment tool.[21] These instruments were selected because they represent different levels of governance. Pakistan provides the national judicial context, UNESCO and CEPEJ offer judiciary-specific international guidance, and the OECD provides a broader public-sector governance perspective.

The documents were read against a common set of questions. First, what responsibility or safeguard does the framework require? Second, what institutional capability would a court need in order to give that requirement practical effect? Third, does the framework indicate how that capability should be assessed before deployment? Finally, what arrangements are required to maintain oversight once an AI system is in use? This approach makes it possible to distinguish between a principle stated in policy and the institutional conditions needed to implement it.

The comparison was conducted thematically rather than through a numerical scoring exercise. Recurring requirements were grouped according to their institutional function. For example, provisions dealing with disclosure, documentation and traceability were considered together as questions of administrative transparency. Requirements concerning responsibility, audit and review informed accountability architecture, while provisions dealing with training, intervention and the ability to question AI outputs were grouped under human oversight capacity. Data quality, privacy, security and system reliability formed a fourth category, digital and data governance capacity. A final category, public legitimacy and engagement, brought together requirements concerning information to users, participation, contestability and confidence in judicial processes. These five dimensions were retained where similar institutional demands appeared across more than one of the frameworks examined.

The study is therefore interpretive rather than an empirical measurement of the readiness of Pakistan's courts. It does not assign readiness scores to individual courts, nor does it claim that the five dimensions exhaust every condition relevant to judicial AI. Instead, the analysis uses existing governance standards to construct a diagnostic framework that can subsequently be tested and refined through institutional application.

The author's professional involvement in Pakistan's recent judicial reform work provides contextual familiarity with the administrative environment in which these policies are being developed. This experience informs the research question but is not treated as primary evidence in the comparative analysis. Findings are based on published policy documents and publicly available institutional material. This distinction is maintained to reduce the risk of treating practitioner knowledge or involvement in the reform process as independent evidence of institutional readiness.

Comparative Findings: From AI Principles to Institutional Requirements

The documents examined in this study differ in purpose and scope. Pakistan's National Guidelines establish a national framework for judicial institutions; UNESCO's guidance is international and judiciary-specific; the OECD approaches AI partly through public-sector governance; and CEPEJ combines ethical principles with a more operational assessment tool. Despite these differences, the comparison shows considerable agreement on the basic safeguards expected around judicial AI. Human control, accountability, transparency, data protection and some form of institutional competence appear repeatedly. The more interesting differences concern how far these principles are translated into conditions that a court can assess before deployment.

Human oversight depends on institutional capacity

Human oversight is the clearest point of convergence. Pakistan's National Guidelines require AI to assist rather than replace judicial decision-making. More specifically, AI-supported processes are subject to human validation, while independent judicial reasoning is retained. UNESCO takes a similar position in its court-specific guidance, treating AI as assistive and requiring meaningful human

[21] National Judicial (Policy Making) Committee, National Guidelines for Use of Artificial Intelligence in Judicial Institutions of Pakistan (Islamabad: Law and Justice Commission of Pakistan, 2026); UNESCO, Recommendation on the Ethics of Artificial Intelligence (Paris: UNESCO, 2021); Juan David Gutiérrez, Guidelines for the Use of AI Systems in Courts and Tribunals (Paris: UNESCO, 2025); OECD, "OECD AI Principles," updated 2024; OECD, Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions (Paris: OECD Publishing, 2025); European Commission for the Efficiency of Justice (CEPEJ).

supervision. CEPEJ's approach is built around user control, and the OECD's examination of justice-sector AI repeatedly returns to human authority over consequential decisions.[22]

There is nevertheless a difference between retaining a human decision-maker and creating meaningful human oversight. The comparison suggests that oversight depends on capabilities surrounding the person who is formally kept "in the loop." A judge or court officer must know when AI has contributed to a process, understand its limitations and have the practical authority to reject an output. Training therefore matters, but training alone is not enough. Institutions also need defined roles, validation procedures and routes for responding when a system behaves unexpectedly.

Examples reviewed by the OECD illustrate the point. Colombia's PretorIA assists the Constitutional Court in identifying cases for review, but judges retain ultimate authority over case selection. In Spain, concerns surrounding the VioGén risk-assessment system have included the extent to which officials follow algorithmic classifications despite having formal discretion to depart from them.[23] Human presence, then, does not establish meaningful oversight by itself. An institution must be able to show how human judgment is protected in practice.

For readiness assessment, the relevant question is therefore not simply whether a human remains involved, but whether that person has the knowledge, authority and procedure necessary to exercise independent judgment.

Accountability has to be located somewhere

Accountability is also common across the frameworks, but it becomes less straightforward once responsibility is translated into an institutional setting. Pakistan's Guidelines place accountability among the principles governing judicial AI. UNESCO likewise insists that responsibility cannot simply be transferred from people or institutions to an AI system, while OECD guidance stresses traceability, risk management and identifiable responsibility.[24]

CEPEJ moves further towards operationalisation. Its Assessment Tool is intended for judicial managers, IT managers and others involved in purchasing, developing or supervising AI systems. It asks who developed the system, the context in which it will operate and the judicial function it will perform. It can also be used in dealings with external providers, allowing information supplied by a vendor to become part of the institution's own assessment.[25] This matters because courts will not necessarily build the systems they use.

The comparison therefore points to what this paper describes as an accountability architecture. Responsibility for judicial AI is rarely located in one person. Procurement may sit with one office, technical management with another, judicial use with judges or court staff, and external development with a private provider. If an AI-supported process produces a serious error, a general statement that "the institution remains accountable" may not tell anyone who must investigate it or who has authority to suspend the system.

Institutional readiness consequently requires responsibilities to be assigned before problems arise. At minimum, a court should be able to identify who authorises a use, who supervises it, who records and reviews incidents, and who can stop or modify deployment.

Transparency is partly an administrative problem

Transparency is often discussed as a property of an AI system: whether an output can be explained, whether the underlying model is sufficiently interpretable, or whether users understand that AI is involved. These concerns are important and appear throughout the documents reviewed. Yet the comparison also points to a less technical form of transparency.

[22] National Judicial (Policy Making) Committee, National Guidelines for Use of Artificial Intelligence in Judicial Institutions of Pakistan (Islamabad: Law and Justice Commission of Pakistan, 2026); Juan David Gutiérrez, Guidelines for the Use of AI Systems in Courts and Tribunals (Paris: UNESCO, 2025); European Commission for the Efficiency of Justice (CEPEJ), European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and Their Environment (Strasbourg: Council of Europe, 2018).

[23] OECD, Governing with Artificial Intelligence, chap. 5, "AI in Justice Administration and Access to Justice." The OECD discusses both Colombia's PretorIA and Spain's VioGén in examining human oversight in justice-sector AI.

[24] National Judicial (Policy Making) Committee, National Guidelines; UNESCO, Recommendation on the Ethics of Artificial Intelligence (Paris: UNESCO, 2021); OECD, Governing with Artificial Intelligence.

[25] CEPEJ, Assessment Tool for the Operationalisation of the European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and Their Environment (Strasbourg: Council of Europe, 2023), 5–8. The tool identifies judicial management bodies and IT/project managers among its intended users and permits information from external providers to inform institutional assessment.

A court needs to know what AI systems are operating within it. It needs records of their purpose, the processes in which they are used and the authority under which they were introduced. Where an AI-generated or AI-assisted output becomes relevant to a judicial process, appropriate documentation may also be necessary for later scrutiny. The OECD's review of AI in justice draws attention to the problem from another direction: evidence about the impact of judicial AI remains fragmented, with relatively few systematic mechanisms for collecting and evaluating outcomes at organisational or national level.[26]

Pakistan's Guidelines provide a useful starting point because transparency is expressly linked with responsible AI use. Pakistan's wider digital reform programme is also developing institutional data capacity. The National Judicial Analytics Dashboard, launched in February 2026, brings together judicial data to support evidence-based policy, performance monitoring and institutional transparency.[27] These developments should not be treated as proof of AI readiness, but they show why existing administrative systems matter. AI governance does not begin with the algorithm. It also depends on whether the institution can document, monitor and explain what it is doing.

Administrative transparency therefore forms a separate readiness dimension in the proposed framework. A court that cannot reliably identify where AI is being used will struggle to audit its effects or explain its role to judges, lawyers and litigants.

Data and digital governance come before sophisticated AI

The strongest connection between technological and institutional readiness appears in data governance. All four approaches recognise, in different forms, that responsible AI depends on reliable data, privacy, security and appropriate technical safeguards. Pakistan's Guidelines expressly require protection of judicial, personal and sensitive information. CEPEJ emphasises data quality and security, while UNESCO and OECD place privacy, robustness and risk management within responsible AI governance.[28]

For courts, however, the issue starts before an AI model is selected. Judicial information may sit across different case-management systems, exist partly in paper records, follow inconsistent formats or contain highly sensitive personal data. A model cannot correct the institutional conditions under which those records were created and maintained. Poor underlying data may instead be reproduced at greater speed.

This is especially relevant in a system undergoing digital transition. Pakistan has moved towards e-filing, judicial analytics and greater automation, but the national reform process itself recognises differences in preparedness across institutions. The decision to initiate district-court e-filing, for example, was accompanied by attention to procedural rules, system preparedness and capacity-building.[29] Such variation means that the readiness of one court or one judicial function cannot simply be assumed for another.

Digital readiness in this framework therefore means more than access to hardware or software. Before AI deployment, an institution should know what data a system will use, whether those data are sufficiently reliable for the proposed purpose, who may access them, how they will be protected and whether the institution can trace their use. This is one area where technological and institutional readiness overlap, but they should not be collapsed into the same concept.

Public legitimacy is recognised, but less consistently operationalised

The fifth dimension is less uniform across the documents. Transparency, human oversight and accountability receive direct treatment in almost every framework. Public participation and institutional legitimacy are more uneven.

UNESCO gives multistakeholder governance a prominent place and connects judicial AI with human rights and public trust. OECD's justice analysis similarly emphasises stakeholder engagement, redress and people-centred justice. CEPEJ places the rights and control of users at the centre of its ethical approach. Pakistan's Guidelines protect due process, fairness and judicial independence,

[26] OECD, Governing with Artificial Intelligence, chap. 5. The report notes that evidence on the impact of justice-sector AI remains fragmented and that systematic organisational and national evaluation mechanisms remain limited.

[27] Law and Justice Commission of Pakistan, "Judiciary Enters the Data Age: National Judicial Analytics Dashboard Launched," February 10, 2026.

[28] National Judicial (Policy Making) Committee, *National Guidelines*; Gutiérrez, *Guidelines for the Use of AI Systems in Courts and Tribunals*; CEPEJ, *European Ethical Charter*; OECD, *Governing with Artificial Intelligence*. Pakistan's Guidelines specifically address protection of judicial, personal and sensitive information.

[29] National Judicial (Policy Making) Committee, "Digital Transformation Roadmap—Launch of E-Filing of Cases," decision recorded by the Law and Justice Commission of Pakistan, 2025–26. The Committee approved e-filing in district courts as the first phase of the national e-court project and required High Courts to address the necessary procedural framework.

although the national framework gives greater attention to institutional safeguards than to a separate mechanism for continuing public participation in AI governance.[30]

For courts, this dimension deserves separate treatment. Judicial AI can affect people who have no choice about whether to interact with the institution using it. A litigant cannot simply select another justice system in the way a consumer might choose another service provider. Nor is public confidence secured merely because an AI system performs accurately. People may reasonably want to know whether AI has materially affected a process, what role it played and how its use can be questioned.

The comparison does not support a claim that every AI application requires the same level of public consultation. A transcription tool and a system influencing case prioritisation present different implications. It does, however, support a broader readiness question: **does the institution have a proportionate way of informing, listening to and providing redress to those affected by the use?**

For this reason, the fifth dimension is described as **public legitimacy and engagement**, rather than public trust alone. Trust cannot simply be required from court users, nor can a court certify itself as trusted. What an institution can assess are the arrangements that may support legitimate use: disclosure, participation where appropriate, accessible explanations and routes through which concerns can be raised.

What the comparison shows

Taken together, the analysis does not reveal a shortage of principles. In fact, the opposite is true. Judicial institutions now have a substantial body of guidance on responsible AI. The recurring difficulty is converting those principles into questions about institutional capability.

The five dimensions identified here perform that conversion. Human oversight becomes a question about training, authority and intervention. Accountability becomes a question about assigned responsibility and review. Transparency requires documentation and institutional visibility over AI use. Data governance requires reliable information systems and controls. Public legitimacy requires disclosure, engagement and routes for challenge appropriate to the risk of the application.

The relationship can be summarised as follows:

Recurring governance principle	Institutional requirement	Readiness question
Human oversight	Knowledge, authority and intervention procedures	Can a responsible human meaningfully scrutinise, reject or override the AI-supported output?
Accountability	Assigned roles, audit and incident responsibility	Is it clear who is responsible for authorising, supervising and responding to failures?
Transparency / explainability	Documentation, disclosure and traceability	Can the court identify and explain where, why and under whose authority AI is being used?
Privacy, security and robustness	Reliable data and digital governance	Are the data and systems sufficiently reliable, secure and governable for the proposed use?
Fairness, user rights and participation	Engagement, information and redress	Can affected users understand the role of AI and raise or challenge concerns where appropriate?

[30] Gutiérrez, *Guidelines for the Use of AI Systems in Courts and Tribunals*; OECD, *Governing with Artificial Intelligence*; CEPEJ, *European Ethical Charter*; National Judicial (Policy Making) Committee, *National Guidelines.* UNESCO's judiciary-specific framework includes multistakeholder governance among its fifteen principles.

The table also clarifies the distinction at the centre of this paper. CEPEJ already provides a serious pre-deployment tool, UNESCO provides operational guidance, and OECD places considerable emphasis on governance mechanisms that should accompany justice-sector AI. The proposed framework should therefore not be presented as the first attempt to assess AI before it reaches a court.

Its contribution lies elsewhere. Existing tools largely begin with an AI system or proposed use and ask whether appropriate safeguards surround it. The framework developed here turns the perspective around. It asks first what the institution receiving that system can demonstrate about its own capacity to govern it.

That difference may appear small, but it changes the sequence of the decision. A technically viable system that satisfies a procurement assessment should not automatically proceed if responsibility inside the court remains unclear, staff cannot exercise meaningful oversight, relevant data are poorly governed, or affected users have no workable route for questioning its use. In those circumstances, the appropriate response is not necessarily to reject AI altogether. It may be to delay deployment, restrict the proposed use or build the missing institutional capacity first.

This is the point at which institutional readiness becomes more than another set of AI principles. It becomes a threshold question for deployment.

An Institutional Readiness Framework for Judicial AI

The comparative analysis suggests that the main difficulty is no longer identifying principles for responsible judicial AI. There is considerable agreement on those principles. The harder question for a court is more practical: what would show that the institution is capable of applying them?

The framework proposed here addresses that question through five dimensions: administrative transparency, accountability architecture, human oversight capacity, digital and data governance capacity, and public legitimacy and engagement. These dimensions are not presented as new ethical principles. Most can already be found, in different forms, in Pakistan's National Guidelines and the international instruments examined earlier. The contribution lies in bringing them together as conditions of institutional readiness and translating them into questions that can be examined before a particular AI use proceeds.

This also changes the sequence of assessment. A court should not begin only with the capabilities of an available AI system and then ask what safeguards need to be added around it. It should first examine the institutional environment into which that system will be introduced. The framework therefore works as a pre-deployment diagnostic, although reassessment would remain necessary after deployment.

Administrative transparency

The first dimension concerns whether a judicial institution has sufficient visibility over its own use of AI. This sounds elementary, but it becomes difficult when tools are introduced at different levels of an organisation or when judges and staff begin using generally available AI applications without a central procurement decision.

At a minimum, the institution should be able to identify which AI systems are being used, for what purposes and under whose authority. It should also distinguish between different kinds of use. An application used to transcribe proceedings raises different issues from one used to summarise pleadings, conduct legal research or assist with case prioritisation. Treating all of them simply as “AI use” provides little basis for meaningful oversight.

Readiness in this area therefore depends on ordinary administrative practices: documentation, approval, record keeping and traceability. Where appropriate to the risk of the application, the institution should also be capable of recording when AI has materially contributed to a process. Transparency begins here, before the more technically difficult question of explaining how a model produced a particular output.

The practical test is whether the court can answer three basic questions without difficulty: Where is AI being used? What is it being used for? Who authorised that use?

If the institution cannot answer them, meaningful oversight is difficult to establish.

Accountability architecture

Knowing where AI is used does not resolve who is responsible for it. Judicial AI may involve several actors: judges, court staff, administrative leadership, IT units, procurement officials and external technology providers. Responsibility can become blurred precisely because no single actor controls the entire process.

The second dimension therefore examines the institution's accountability architecture. Before deployment, responsibility should be identifiable at different stages of the system's use. Someone must authorise the application. Someone must be responsible for its technical administration. There should be a clear route for reporting errors or unintended consequences, and authority to restrict or suspend its use where necessary.

This does not mean transferring responsibility for judicial decisions from judges to administrators or technology providers. Judicial independence must remain intact. The point is narrower: independence should not produce an accountability gap around the technology supporting judicial work.

A court should consequently be able to identify responsibility for at least four functions: authorisation, supervision, incident response and review. The precise institutional actor performing each function may differ between courts and applications. What matters for readiness is that the functions have been assigned rather than assumed.

Human oversight capacity

Human oversight appears in virtually every major framework considered in this study. Its practical meaning, however, depends heavily on institutional capacity.

Keeping a human formally involved in a process is not sufficient if that person routinely accepts an AI-generated output without understanding its limitations or lacks the authority to challenge it. Meaningful oversight requires both competence and room for independent judgment.

Three issues are particularly relevant before deployment. The first is knowledge. Judges or staff using an application should understand its intended function and its important limitations. They do not need to become computer scientists, but they should know enough to recognise circumstances in which reliance on an output would be inappropriate.

The second is authority. Users must be able to disregard, correct or escalate an AI output without being institutionally pressured to accept it merely because the system produced it. The third is procedure. The court needs some mechanism for dealing with errors that are not isolated. If several users identify the same problem, that information should reach whoever has authority over continued deployment.

This dimension therefore asks something more demanding than whether a “human is in the loop.” It asks whether the human can exercise judgment inside that loop.

Digital and data governance capacity

AI does not enter an empty technological environment. It relies on existing records, databases, case-management systems, security arrangements and data practices. Their condition affects what an AI application can reasonably be expected to do.

This dimension is particularly important for judicial systems undergoing digital transition. The existence of digitised records should not automatically be interpreted as evidence that those records are suitable for AI use. Data may be incomplete, inconsistently classified, duplicated or collected under different administrative practices. Sensitive judicial information also raises questions about access, storage, retention and the involvement of external providers.

A pre-deployment assessment should therefore examine the data required by the proposed application rather than asking generally whether the court is “digitised.” The institution should know where the data come from, whether their quality is adequate for the intended function, who can access them and how their use can be traced. Security and privacy arrangements should correspond to the sensitivity of the information involved.

This is where technological and institutional readiness meet most directly. A technical team may determine whether a dataset can be used by a model. The institutional question is whether it should be used for that purpose and whether the court can govern its use responsibly.

Public legitimacy and engagement

The final dimension concerns the people who encounter the justice system. Unlike the preceding four dimensions, it cannot be reduced to an internal organisational process.

Not every use of judicial AI requires public consultation. Requiring the same engagement process for automated transcription and an application that influences case prioritisation would make the framework unnecessarily rigid. The level of disclosure and engagement should correspond to the significance of the AI system within the judicial process.

There is nevertheless a minimum institutional question. Where AI can materially affect a person's interaction with the justice system, does that person have a reasonable way of understanding its role and raising a concern about its use?

This may require different mechanisms in different circumstances: public information about approved applications, notification in higher-risk uses, accessible explanations, complaint procedures or an opportunity to challenge an AI-supported process. The framework does not assume that these arrangements will automatically produce public trust. Trust is not something an institution can declare or require. What courts can do is establish procedures that make the use of AI more visible, contestable and consistent with procedural fairness.

For that reason, public legitimacy, rather than public trust alone, is the appropriate readiness criterion.

Moving from dimensions to a readiness decision

The five dimensions are intended to lead to a decision, not simply another checklist. This requires some indication of what different levels of readiness mean.

A three-level classification is proposed:

Readiness level	Institutional position	Appropriate response
Level 1: Not Ready	Important arrangements are absent, unclear or largely informal. Responsibility may be uncertain, users inadequately prepared, or relevant data and oversight arrangements unreliable.	Do not proceed with deployment. Address the identified institutional gaps first.
Level 2: Developing	Relevant policies and capacities exist but are incomplete, unevenly implemented or insufficiently tested.	Consider a limited pilot only where the use is proportionate to the identified risk and additional safeguards are in place.
Level 3: Ready for Controlled Deployment	Responsibilities are assigned, users are prepared, relevant data are governed, oversight procedures operate in practice, and proportionate transparency and redress mechanisms exist.	Deployment may proceed subject to monitoring and periodic reassessment.

The term controlled deployment is deliberate. Level 3 should not be understood as permanent certification that an institution is ready for any form of AI. Readiness is tied to a particular application and its likely consequences.

This matters because judicial AI varies considerably in risk. Consider two hypothetical uses in the same court. The institution may have reliable procedures for an AI transcription service: staff can verify transcripts, mistakes can be corrected and the original audio remains available. The same court might consider an AI tool that recommends which cases should receive priority. That second application could require stronger data controls, clearer explanations, more extensive testing and greater procedural safeguards. The institution's general technological capacity has not changed, but the readiness decision should.

The framework therefore rejects a single institution-wide label of “AI-ready.” Assessment takes place at the intersection of institutional capacity and proposed use.

The proposed diagnostic

The complete framework can be expressed through five questions:

Dimension	Core readiness question	Examples of evidence
Administrative transparency	Does the institution know where, why and under whose authority AI will be used?	AI register; documented purpose; approval record; use protocols
Accountability architecture	Are responsibilities for the AI-supported process clearly assigned?	Named responsible units; audit arrangements; incident procedures; suspension authority
Human oversight capacity	Can relevant personnel meaningfully scrutinise and intervene in AI-supported processes?	Training records; validation procedures; override authority; escalation mechanisms
Digital and data governance	Are the data and digital systems supporting the application reliable and governable for its intended purpose?	Data-quality assessment; access controls; privacy safeguards; security procedures; audit trails
Public legitimacy and engagement	Are people affected by the application appropriately informed and able to question its use?	Disclosure arrangements; accessible information; consultation where warranted; complaint or review mechanisms

The evidence column is illustrative rather than mandatory. A small administrative application should not be required to satisfy the same documentation burden as an AI system used close to adjudication. The framework is intended to impose discipline on the decision, not bureaucracy for its own sake.

One further safeguard is necessary. The five dimensions should not simply be averaged into a numerical score. A court performing strongly on four dimensions should not be able to compensate mathematically for a serious failure in the fifth. Strong digital infrastructure, for example, cannot cancel out the absence of meaningful human oversight. For higher-risk applications, some conditions should operate as minimum thresholds rather than weighted preferences.

This is where the framework differs from a conventional technology-readiness index. Its purpose is not to produce a league table of courts. It is to identify whether the institutional conditions required for a proposed AI use are present, which conditions remain weak, and what should happen before deployment proceeds.

The sequence can therefore be stated simply:

Proposed AI use → institutional readiness assessment → address identified gaps → controlled deployment or pilot → continuing review.

Under this approach, the decision to delay deployment is not evidence that a judicial institution has failed at digital transformation. It can instead be evidence that governance is working. A court that recognises an accountability, data or oversight problem before introducing a system is in a stronger position than one that discovers the same problem after the technology has begun affecting judicial processes.

The framework ultimately changes what is meant by being “ready” for judicial AI. Readiness is not demonstrated by possessing sophisticated technology. It is demonstrated by an institution's ability to remain responsible for what that technology does within the justice system.

Applying Institutional Readiness to Judicial Reform in Pakistan

The framework developed in the previous section is not intended to produce a general verdict on whether Pakistan's judiciary is “ready” or “not ready” for artificial intelligence. Such a conclusion would require evidence from individual courts on their data systems, staff capacity, internal procedures and actual use of AI. The available policy material does not provide that evidence. What it does allow is a more limited assessment of where Pakistan's current reform architecture already supports institutional readiness, and where further operational work would be needed before particular AI applications are deployed.

Pakistan is not starting from zero. The National Guidelines provide a governance baseline, digitalisation is progressing through initiatives such as e-filing and judicial analytics, and the institutional architecture for judicial automation has already been

established through the National Judicial Automation Committee.[31] The question is therefore less about creating another national statement of principles and more about translating the principles already adopted into decisions that can be made at court level.

From national guidelines to institutional implementation

One strength of Pakistan's approach is that the National Guidelines do not treat AI as a replacement for judicial judgment. Human oversight, accountability, privacy, fairness and capacity-building are already recognised.[32] The Guidelines also allow the High Courts to develop implementation arrangements suited to their own institutional circumstances. That flexibility is sensible in a judicial system where administrative capacity and digital infrastructure are unlikely to be identical across jurisdictions.

It also means that implementation becomes particularly important. A national commitment to human oversight, for example, still leaves practical questions for an individual institution. Which applications require prior approval? Who may use them? What information may be entered into external generative-AI systems? When must an AI-generated output be independently checked? Who should be informed when a system produces a serious or recurring error?

The readiness framework can help translate a general principle into these institutional questions. It does not require the national framework to prescribe identical procedures for every court. Instead, it provides a common basis against which different implementation arrangements can be assessed.

Begin with an inventory of actual AI use

The first practical step is relatively modest. Judicial institutions need an accurate picture of where AI is already being used.

Formal procurement records will not necessarily provide that picture. Generative AI differs from earlier court technologies because individual judges and staff can access publicly available tools without waiting for an institution to acquire a dedicated system. UNESCO's 2026 consultations in Pakistan are significant in this respect: participants reported uses relating to legal research, drafting and case management alongside concerns about training and institutional preparedness.[33]

A readiness assessment should therefore begin with an institutional inventory covering both formally adopted applications and, as far as practicable, authorised categories of externally available tools. The purpose would not be to police every technological interaction. It would establish where AI is entering judicial work, what information it may process and which uses deserve closer scrutiny.

This would also make later governance possible. A court cannot supervise a use that it does not know exists.

Different uses should face different thresholds

Pakistan should also avoid treating judicial AI as a single category of risk. The appropriate institutional threshold should depend partly on how close an application comes to the exercise of judicial authority.

A useful starting distinction could be made between three broad types of use:

Type of use	Illustrative functions	Readiness implication
Administrative support	transcription, scheduling, document organisation	Basic governance requirements; human verification and data safeguards
Professional support	legal research, summarisation, drafting assistance	Stronger requirements for verification, confidentiality, disclosure and user competence
Decision-proximate use	case prioritisation, risk classification, recommendations capable of influencing judicial reasoning	Highest threshold for oversight, accountability, testing, explanation and review

[31] Supreme Court of Pakistan, "National Judicial Automation Committee Resolves to Push Digital Transformation of Judiciary," May 7, 2025; National Judicial (Policy Making) Committee, "Digital Transformation Roadmap—Launch of E-Filing of Cases," 2025–26.
[32] National Judicial (Policy Making) Committee, National Guidelines for Use of Artificial Intelligence in Judicial Institutions of Pakistan (Islamabad: Law and Justice Commission of Pakistan, 2026).
[33] UNESCO, "AI and Courts: A Tale of Two Intelligent Entities," July 14, 2026.

These categories are illustrative rather than fixed. A seemingly administrative application can become consequential depending on how it is used. Case prioritisation, for example, may appear administrative but could materially affect how quickly a litigant receives judicial attention. Classification should therefore depend on the effect of the application rather than its label.

The advantage of this approach is that it avoids two extremes. Pakistan need not subject every low-risk automation tool to an elaborate approval process, but neither should a technically functional application move closer to adjudication merely because a human remains formally responsible for the final decision.

Readiness should be demonstrated, not assumed

Pakistan's continuing digitalisation provides some of the infrastructure on which future AI applications may depend. E-filing and the National Judicial Analytics Dashboard, for example, demonstrate an institutional move towards structured digital information and data-informed judicial administration.[34] They should not, however, be treated as proxies for AI readiness.

For a proposed AI application, an institution would still need to examine the particular data involved. Are the records sufficiently complete and consistent for this purpose? Do they contain confidential or personally sensitive material? Will information leave judicial infrastructure or be processed by an external provider? Can access be traced? How will errors in the underlying data be corrected?

The same principle applies to staff capacity. The existence of a training programme is not by itself evidence of meaningful human oversight. The relevant evidence is whether the people expected to supervise a particular application understand its function and limitations and know what to do when its output appears unreliable.

This is why the framework emphasises **demonstrated capacity**. Policies, training plans and technical infrastructure are relevant evidence, but readiness ultimately depends on whether the arrangements work for the proposed use.

Institutional ownership needs to be clear

Implementation will also require clarity about institutional ownership. Pakistan's judicial structure makes this especially important. National policy coordination, High Court autonomy, judicial independence and the administrative responsibilities of registries and technical units all need to coexist.

A workable model would not require a single body to control every aspect of judicial AI. It would require responsibility to be visible. For each approved application, the relevant institution should be able to identify who authorised its use, who is responsible for technical management, who monitors incidents and who has authority to suspend or reconsider deployment. Where an external provider is involved, contractual responsibility should not replace institutional responsibility.

The existing National Judicial Automation Committee provides a potential coordination setting at the national level, while implementation can remain within the constitutional and administrative authority of the relevant courts. This would allow common standards and institutional learning without requiring every court to adopt the same technological solution.

Public legitimacy should enter before controversy does

Public communication about judicial technology often occurs after something has gone wrong. For AI, that would be a poor governance strategy.

Not every application requires extensive consultation, but uses capable of materially affecting litigants should have proportionate arrangements for disclosure and challenge. At minimum, people should be able to understand when AI has played a significant role in a process affecting them and where concerns about that use can be raised. For higher-risk applications, consultation with the legal profession, court users and relevant technical or civil-society expertise may be appropriate before full deployment.

This is particularly important because public trust should not be treated as an automatic benefit of technological modernisation. Introducing AI into a court does not make the institution more trustworthy simply because the technology is advanced. Nor does public scepticism necessarily indicate resistance to innovation. In a judicial setting, asking how a consequential process works and who remains responsible for it is a legitimate expectation.

[34] Law and Justice Commission of Pakistan, "Judiciary Enters the Data Age: National Judicial Analytics Dashboard Launched," February 10, 2026; National Judicial (Policy Making) Committee, "Digital Transformation Roadmap—Launch of E-Filing of Cases."

The objective should therefore be **institutional legitimacy rather than technological acceptance**.

A practical pathway for Pakistan

Taken together, these considerations suggest a relatively simple implementation sequence for Pakistan:

Identify the proposed use → classify its institutional significance → assess the five readiness dimensions → document unresolved gaps → pilot where appropriate → review evidence from the pilot → approve, modify, suspend or scale the application.

This process could be incorporated into implementation arrangements developed under the National Guidelines rather than requiring a new national AI policy. For lower-risk applications, the assessment could remain brief. Applications closer to adjudication would require stronger evidence across all five dimensions.

A pilot should also have an exit route. Institutions sometimes treat a pilot as the first stage of inevitable expansion. Under a readiness approach, it is an evidence-generating stage. If the pilot reveals weak data, persistent errors, inadequate human scrutiny or effects on procedural fairness that cannot be satisfactorily addressed, discontinuing the application should remain a legitimate outcome.

The same framework can support gradual adoption. An institution classified as *Developing* need not wait until every aspect of its administration is technologically sophisticated. It can identify the particular weaknesses relevant to the proposed use, address them and reassess. Readiness therefore becomes part of institutional development rather than a barrier to innovation.

Pakistan's present position makes such an approach feasible. The country already has a national judicial AI framework, an institutional mechanism for judicial automation and an expanding digital reform programme. What is needed next is not another catalogue of AI ethics. It is a way of determining whether those principles can be carried through in the institutions where AI will actually be used.

The practical implication is modest but consequential: before asking what AI can do for a court, the institution should establish what it must be able to do to govern that AI.

Conclusion

The debate around artificial intelligence in courts has moved quickly. The question is no longer simply whether AI has a place in judicial administration. Courts are already experimenting with it, international standards have become more detailed, and Pakistan now has its own national guidelines for judicial AI. The more immediate issue is how judicial institutions decide that they are prepared to use these technologies responsibly.

This paper has argued that technological readiness provides only part of that answer. A court may have digitised records, technical infrastructure and access to capable AI tools without having equally developed arrangements for accountability, human oversight, data governance or public-facing transparency. These institutional conditions become especially important in the judiciary because the consequences of technological failure are not confined to administrative inefficiency. They may affect procedural fairness, judicial independence and confidence in the administration of justice.

The comparative analysis found substantial agreement across Pakistan's National Guidelines and the UNESCO, OECD and CEPEJ frameworks. The major principles are already well established. The problem is less one of identifying additional principles than of determining whether courts have the capacity to put existing ones into practice. The Institutional Readiness Framework developed in this paper brings these requirements together under five dimensions: administrative transparency, accountability architecture, human oversight capacity, digital and data governance capacity, and public legitimacy and engagement. It then treats readiness as a use-specific decision, ranging from *not ready* to *developing* and *ready for controlled deployment*.

For Pakistan, this approach would complement rather than replace the governance architecture now being developed. The National Guidelines provide the normative foundation. An institutional readiness assessment could help translate that foundation into decisions about particular AI applications at court level, while leaving room for differences in institutional capacity and the level of risk involved.

There are limits to what can be concluded from this study. The framework has been derived from comparative policy analysis and has not yet been empirically tested across Pakistani courts. Nor does the analysis attempt to measure current levels of readiness within individual judicial institutions. A useful next stage of research would therefore apply the framework through interviews with judges, court administrators and technical personnel, institutional case studies, and pilot assessments of particular AI applications.

Such work could test whether the five dimensions capture the practical constraints courts encounter and whether the proposed readiness levels require modification.

For now, the paper makes a more limited claim. The responsible adoption of judicial AI should not begin with the assumption that technological capability establishes preparedness. Before deciding whether an AI system is suitable for a court, there is an earlier question to answer: is the court institutionally equipped to remain accountable for its use? For judicial reform, that question should come first.

Bibliography

1. Barysė, Dovilė, and Roee Sarel. "Algorithms in the Court: Does It Matter Which Part of the Judicial Decision-Making Is Automated?" *Artificial Intelligence and Law* 32 (2024): 117–146. https://doi.org/10.1007/s10506-022-09343-6.
2. Dancy, Tatiana, and Monika Zalnieriute. "AI and Transparency in Judicial Decision Making." *Oxford Journal of Legal Studies* 46, no. 1 (2026): 1–34. https://doi.org/10.1093/ojls/gqaf030.
3. European Commission for the Efficiency of Justice (CEPEJ). *Assessment Tool for the Operationalisation of the European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and Their Environment*. Strasbourg: Council of Europe, 2023.
4. European Commission for the Efficiency of Justice (CEPEJ). *European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and Their Environment*. Strasbourg: Council of Europe, 2018.
5. Gutiérrez, Juan David. *Guidelines for the Use of AI Systems in Courts and Tribunals*. Paris: UNESCO, 2025.
6. Law and Justice Commission of Pakistan. "Judiciary Enters the Data Age: National Judicial Analytics Dashboard Launched." February 10, 2026.
7. Legara, Erika Fille, Elmo Domino Jose, and Paula Joy Martinez. "Beyond Model Readiness: Institutional Readiness for AI Deployment in Public Systems." Paper presented at the 2nd Workshop on Technical AI Governance Research, ICML 2026, 2026.
8. National Judicial (Policy Making) Committee. "Digital Transformation Roadmap—Launch of E-Filing of Cases." Decision recorded by the Law and Justice Commission of Pakistan, 2025–26.
9. National Judicial (Policy Making) Committee. *National Guidelines for Use of Artificial Intelligence in Judicial Institutions of Pakistan*. Islamabad: Law and Justice Commission of Pakistan, 2026.
10. OECD. *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*. Paris: OECD Publishing, 2025. https://doi.org/10.1787/795de142-en.
11. OECD. "OECD AI Principles." Updated 2024.
12. Reiling, A. D. (Dory). "Courts and Artificial Intelligence." *International Journal for Court Administration* 11, no. 2 (2020): 1–10. https://doi.org/10.36745/ijca.343.
13. Supreme Court of Pakistan. "National Judicial Automation Committee Resolves to Push Digital Transformation of Judiciary." Press release, May 7, 2025.
14. Susskind, Richard. *Online Courts and the Future of Justice*. Oxford: Oxford University Press, 2019.
15. Tyler, Tom R. "Procedural Justice and the Courts." *Court Review* 44, nos. 1–2 (2007–08): 26–31.
16. UNESCO. "AI and Courts: A Tale of Two Intelligent Entities." July 14, 2026.
17. UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO, 2021.
18. Zalnieriute, Monika, Lyria Bennett Moses, and George Williams. "The Rule of Law and Automation of Government Decision-Making." *Modern Law Review* 82, no. 3 (2019): 425–455. https://doi.org/10.1111/1468-2230.12412



2026 by the authors; Journal of Global Social Transformation (JGST). This is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) license (http://creativecommons.org/licenses/by/4.0/).