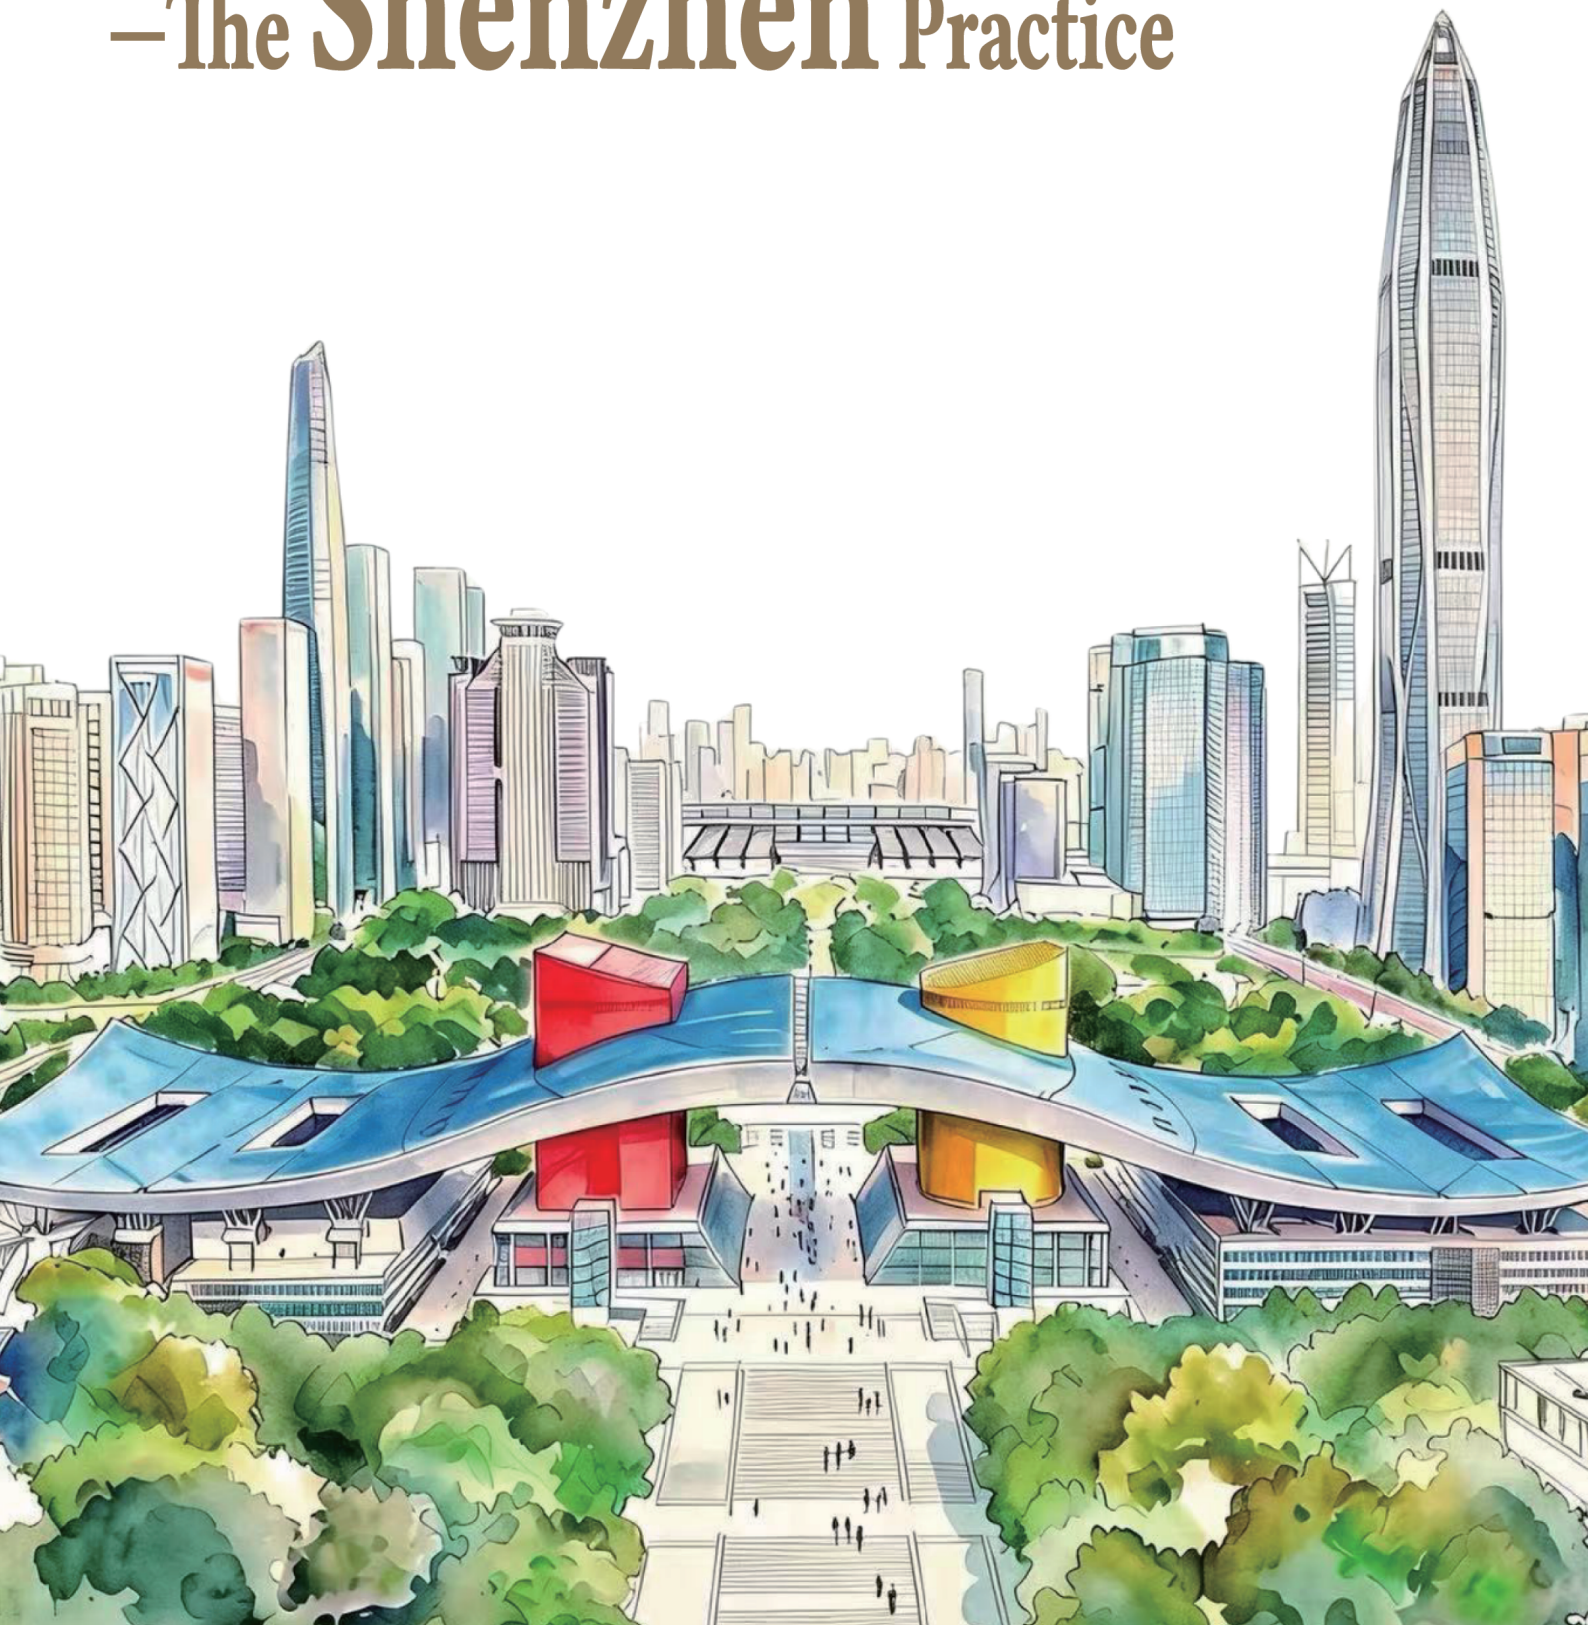
AI in the Judiciary

—The Shenzhen Practice



AI in the Judiciary

–The Shenzhen Practice

Political and Legal Affairs Commission of the CPC Shenzhen Municipal Committee

Shenzhen Intermediate People's Court

Government Services and Data Management Bureau of Shenzhen Municipality

September 2026



未来已来，智在人为

The Future is Present， Wisdom Driven by Actions

Table of Contents

Foreword I 02
Foreword II 04
Preface 07

1 Background to the Shenzhen Practice 09
1.1 / The Challenge: Efficiently Resolving Disputes under Heavy Caseloads 09
1.2 / The Opportunity: A Technical Breakthrough and a Policy Framework 10
1.3 / The Foundation: Paperless Case Handling, Online Litigation Services, and Knowledge Management 11

2 Systemic Outcomes of the Shenzhen Practice 14
2.1 / The System's Six Core Application Scenarios 14
2.2 / Agents in the System 23
2.3 / A System That Grows 27

3 Innovations of the Shenzhen Practice 29
3.1 / The First Systematic, End-to-End Judicial AI Application 29
3.2 / The First Domain-Specific Harness Engineering for the Judiciary 30
3.3 / The First Judicial AI Security Framework Deployed in Production 34

4 Practical Application Outcomes of the Shenzhen Practice 36

5 Experience from the Shenzhen Practice 39
5.1 / Concentrating Resources: Higher-Court Guidance and Local Support 42
5.2 / Rules First: Setting the Direction and Holding the Line 43
5.3 / Adjudication-Led, Cross-Disciplinary Development 44
5.4 / Growing the Ecosystem:Changing Mindsets and Deepening Application 46

6 A Vision for Digital-Intelligence-Native Justice 49
6.1 / Long-Term Challenges and Contemporary Opportunities in Global Judicial Work 49
6.2 / Creating a "Digital-Intelligence-Native" Mode of Judicial Work 52
6.3 / Building a Smart and Convenient Online Litigation Service System 56
6.4 / Expanding the Public Service and Social Governance Functions of Judicial AI 60

Conclusion 62
Appendix 1 Guide for Using AI Agents in the Judicial Domain V2.0 (Sept 2026) 63
Appendix 2 Glossary 75
Appendix 3 Chronology of Major Events 79

Foreword I

Xi Jinping Thought on the Rule of Law Guides the Judicial Application of Artificial Intelligence

A new round of technological and industrial transformation is gathering pace. Artificial intelligence (AI) is profoundly reshaping how people work and live, while profoundly transforming the development of the rule of law and social governance. General Secretary Xi Jinping has emphasized the need to respect the principles governing judicial work, combine deeper judicial reform with the application of modern technology, and promote "the deep integration of technological innovations such as big data and AI with judicial work." In accordance with the decisions and arrangements of the CPC Central Committee and of the Provincial and Municipal Party Committees, Shenzhen's political and legal organs have thoroughly implemented Xi Jinping Thought on the Rule of Law and the principles of General Secretary Xi Jinping's important speeches and instructions, strengthened the intelligent development of political and legal work, and explored new approaches to law enforcement and judicial work that are deeply integrated with modern technology. Shenzhen Courts have built the world's first end-to-end AI-Assisted Adjudication System and pioneered a framework of rules for the judicial application of AI, offering a practical model for strengthening judicial capacity under resource constraints.

I. Easing the tension between rising caseloads and limited judicial resources.

Persistently rising caseloads, constrained judicial resources, and mounting pressure on case handling have become a structural problem and a practical challenge for modern judicial systems. Shenzhen is a major economic and foreign-trade hub as well as a megacity: caseloads are heavy and judicial needs are diverse, and the traditional model of adjudication struggles to meet the public's expectations of fairness and justice. With a mandate to pioneer and demonstrate, Shenzhen Courts seized the opportunity created by AI capacity breakthroughs and turned to large language model (LLM) powered AI to address this challenge, developing the Shenzhen Courts' AI-Assisted Adjudication System (hereafter the System). The System covers the key stages across the entire case-handling process–case filing, case-file review,

hearings, and judicial-document drafting–and brings together AI agent assistance, data support, process supervision, and continuous self-improvement. It has improved the quality and efficiency of adjudication, eased the imbalance between judicial workload and staffing, and advanced reform of adjudication models. The System's functions have now been incorporated into the unified national digital platform for case handling and court administration for People's Courts Nationwide, operated by the Supreme People's Court of the People's Republic of China (SPC), and are being tried in courts across China. The accompanying framework of rules for the judicial application of AI has gained wide recognition from judicial communities in China and abroad.

II. Implementing the judicial accountability system fully and accurately.

Judicial power is, in essence, the power to make judgments, and that power must remain firmly in human hands. However far technology advances, a judge's professional judgment and judicial experience are irreplaceable: adjudicative authority belongs exclusively to judges, who bear the ultimate judicial responsibility. The System strictly positions AI only as a supporting tool. Judges retain complete control over the System's information sources, decision-making processes, and outputs, thereby preventing at the source the risks of "Machine Adjudication" and an "algorithmic black box". Authoritative data such as statutory provisions and reference cases come entirely from the SPC knowledge base rather than being generated directly by LLMs, which avoids the risk of information distortion caused by hallucinations. The System standardizes case-handling and decision-making workflows: at each stage, it provides intelligent retrieval, automatic risk alerts, and analytical support, while every point where risk arises can be traced and reviewed. Security is treated as the System's lifeline–it is physically isolated from the internet, while data flows are monitored in real time and subject to various types of risk alerts. In the emerging field of judicial AI, the System offers a compelling example of the judicial accountability system being implemented fully and accurately.

III. Advancing technology-enabled modernization of social governance.

As a Pilot Demonstration Area of Socialism with Chinese Characteristics and a national innovative city, Shenzhen is carrying forward the "AI+" initiative, integrating technological innovation with political and legal work and applying AI, internet-based technologies, and big data to enhance urban governance capacity and advance its modernization. From AI-enabled detection of disputes and early risk warnings to refined urban governance, technology has become an important means of improving governance. The System, with its large-scale real-world application, provides a compelling real-world example of technology-enabled governance. Shenzhen will continue to develop and apply modern technology across a wide range of social-governance settings, translating technological advantages more effectively into better governance outcomes.

Upholding Fundamental Principles while Pursuing Innovation. Steadfast action carries us farther. The white paper AI in the Judiciary–The Shenzhen Practice embodies Shenzhen's efforts to integrate modern technology with judicial practice and demonstrates its confidence and resolve to make justice fairer and more efficient. We look forward to working together with more partners and contributing greater insight and energy to the advancement of justice and broader social progress.

With these words, we offer this preface.

Political and Legal Affairs Commission of the CPC Shenzhen Municipal Committee

August 2026

Foreword II

Practical Insights into the Application of Artificial Intelligence in Government

AI, a key driving force behind a new round of technological and industrial transformation, is reshaping the global development landscape with unprecedented depth and reach. When presiding over the ninth group study session of the Political Bureau of the 19th CPC Central Committee in October 2018, General Secretary Xi Jinping used the word "strategic" four times to underline the importance of artificial intelligence: AI is "a strategic issue bearing on whether China can seize the opportunities of this new round of technological and industrial revolution"; "a strategic technology driving this round of technological revolution and industrial transformation"; "an important strategic instrument for gaining the initiative in global technological competition"; and "a key strategic resource for driving China's leap in science and technology, the upgrading of its industries and an overall rise in productivity." These remarks illuminate the strategic position and catalytic role of AI in economic and social development and chart a course for its application in government.

Since the National Data Administration was established in October 2023, its work has centred on market-oriented allocation of data elements, coordinating the development of Digital China, the digital economy, and a digital society, as well as international cooperation in the data sector. It has developed a "5+3+1" framework for reforming the market-oriented allocation of data elements, with the "1" representing one key priority: using data elements to empower innovation and development in AI. In a word, artificial intelligence and data elements have become the two core engines driving development in the digital age.

Shenzhen is a national centre of the electronic information industry: the global competitiveness of its new-generation electronic information industry continues to strengthen, and its core AI industry ranks in the country's first tier by scale. As the wave of AI gathered momentum, Shenzhen had strong foundations to draw on, but also faced real difficulties in government settings–unclear implementation paths, absence of unified standards, and difficulty in measuring computing-power capacity. For the agencies responsible for

coordinating government services and data, how to make new technologies deliver real value in the digital transformation of government, especially to advance innovations in service delivery and governance approaches while providing more real-world application scenarios for industrial development, became a defining question that needs an answer.

Since the release of ChatGPT in November 2022, it has drawn worldwide attention, demonstrating the unprecedented generalization and content-generation capabilities of LLMs. The technology was then at a very early stage: there was neither a mature path for deploying it extensively in public service nor a proven implementation precedent to follow. There were no ready answers to how it should be used or what it could achieve. At that point, Shenzhen Courts took the lead, charting a successful path for AI-assisted adjudication. Recognizing the value of that work, we established a standing liaison mechanism and provided full support for project-plan revisions, computing services, and data coordination. By March 2024, the pilot had begun to deliver results: in industrial-design patent cases, the same caseload was handled by a team reduced from twenty members to eight, with the number of judges falling from eight to two—an encouraging outcome. On 28 June 2024, the System was rolled out to civil and commercial cases citywide, marking the transition of AI-assisted adjudication from pilot testing to full-scale application. It gradually came to cover the civil and commercial cases that make up the great majority of Shenzhen Courts' caseload, with several thousand judges and judicial support staff using it simultaneously. Real use at that scale created a feedback flywheel, and model quality rose quickly. By the end of December that year, the System had been extended to administrative and criminal cases, achieving wide coverage of case types and providing a "Shenzhen model" for deploying LLMs in public service departments across China.

As successive waves of technological change have surged forward, Shenzhen Courts have never stopped exploring. We have tracked that evolution and maintained close, regular coordination with Shenzhen Courts, working together on the data governance, computing-capacity expansion, and model tuning that each round of iteration and upgrading required. After DeepSeek was released as open source, Shenzhen Courts were the first to begin adapting the System to Huawei's Ascend compute platform; and when the "AI Agents+Skills" pattern of Agent application emerged at the end of 2025, Shenzhen Courts went on to explore a new judicial paradigm in which AI agents work in concert with the Judicial Domain-Specific Harness Engineering. From closed-source experiment to open-source adoption and then to the spread of Agents, Shenzhen Courts have stood at the front line of innovation and application through every wave.

Looking back, Shenzhen Courts' work stands out for three distinct qualities. First, a pioneering spirit: Shenzhen Courts dared to break new ground and achieved tangible results in a focused application scenario within six months. Second, a systematic, needs-driven approach: with fairness and efficiency as the central goal, Shenzhen Courts built an end-to-end method combining operational leadership, organizational safeguards, and technical support, establishing a replicable and scalable path to implementation. Third, a capacity for continuous iteration: the technology changes, but the focus on using new quality productive forces—that is, productivity driven by technological innovation and other advanced productive capabilities—to solve real governance problems remains constant, keeping the initiative at the forefront of the field. In essence, Shenzhen Courts have used new quality productive forces to resolve governance problems that had long resisted conventional means. This is not simply replacing people with machines; it is reshaping the mode of judicial work through human-AI collaboration, so that judges can concentrate on the core functions of adjudication.

A successful pilot can pave the way for broader adoption. From May 2024, we have promoted the Shenzhen Courts model across the city through on-site visits, rollout meetings, and the drafting of standards and specifications. AI applications across Shenzhen's government have moved from isolated success to widespread adoption, and now deliver results in government services, public concerns and requests, traffic management, city management, environmental protection, and other fields.

This white paper systematically reviews Shenzhen Courts' three-year journey of exploration and provides a comprehensive account of Shenzhen's experience in taking LLM applications in government from initial deployment to large-scale adoption. We hope that this first-hand account will offer ideas and methods that other cities and government departments can draw on, allowing the synergies between AI and data elements to be fully realized across an even broader range of governance settings.

Government Services and Data Management Bureau of Shenzhen Municipality
August 2026

The Shenzhen Practice of Artificial Intelligence in the Judiciary

Preface

The world is undergoing changes of a scale not seen in a century, together with a new round of technological and industrial revolution. AI has become one of its principal drivers–a new engine of global growth and an accelerator in the shift from old growth drivers to new ones. It is reshaping how people produce and live, how societies are governed, and how the rule of law itself is built.

In January 2019, General Secretary Xi Jinping addressed the Central Political and Legal Affairs Conference and called for "the deep integration of technological innovations such as big data and AI with judicial work." The Opinions of the CPC Central Committee on Strengthening Adjudication Work in the New Era call for wider digital and intelligent application in the courts, and for these technologies to drive change in adjudication procedures, litigation rules, and judicial models. The Sixth Five-Year Reform Outline of the People's Courts (2024–28) goes further: it calls for indigenous innovation in judicial technology; for greater use of intelligent algorithms, LLMs, and digital models; for AI to assist case handling, consolidate routine work, control risk, and ease litigation; and for a judicial technology system that is secure and reliable, efficient and coordinated, accessible and convenient to all–driving that same transformation of adjudication procedures, litigation rules, and judicial models.

Shenzhen is a city built on innovation. It is dense with technology companies, ranks among the country's leaders in the high-tech industry, has a highly active market economy, and generates new business forms and models continuously. Judicial demand here is correspondingly large and varied, and two pressures weigh especially heavily on Shenzhen Courts: high caseloads and the difficulty of resolving disputes conclusively. Incorporating AI into judicial work offers a way through.

In June 2024, Shenzhen Courts completed the world's first systematic AI-Assisted Adjudication System. Extended step by step, it now covers civil, administrative and criminal proceedings, supports judges across the whole case-handling process, and is now deeply embedded in day-to-day judicial work. The system has improved both the fairness and the efficiency of adjudication, and it marks the step from blueprint

to working practice for judicial AI. To match this new mode of working, Shenzhen Courts were also the first to build a mechanism for specialized adjudication of categorized case types, and to chart a way of fully and accurately implementing the judicial accountability system under human-AI collaboration–a practical model for integrating AI with social governance more broadly. The rules for the judicial use of AI, developed while building and running the system, have yielded practices that other courts can replicate, and are now widely recognized by the judicial community in China and abroad.

Three years of practice have shown that AI can be used safely in the judicial field, and that it can drive real change in adjudication procedures, litigation rules, and the way justice is delivered. Deep integration of AI with adjudication puts Xi Jinping Thought on the Rule of Law into practice and offers a route through high caseloads and limited judicial staffing to higher quality and efficiency of adjudication. Ultimately, it underpins the modernization of governance in a megacity, the modernization of adjudication, and the building of a socialist judicial system that is fair, efficient, and credible.

This white paper sets out what Shenzhen Courts have learned: the innovations, the institutional rules that emerged, and the trends now visible, so that the experience can be shared with courts deepening their own digital construction, and with all looking to the modernization of adjudication.



1. Background to the Shenzhen Practice

1.1 The Challenge: Efficiently Resolving Disputes under Heavy Caseloads

> **"The judiciary is the last line of defense to safeguard social fairness and justice."**
>
> **– Xi Jinping**

Experience around the world shows that the volume of disputes tends to rise with the pace and level of economic development. In just a few decades, China compressed an industrialization process that took developed countries centuries, achieving rapid economic growth while maintaining long-term social stability. As the economy expanded, society transformed and urbanization accelerated, and disputes multiplied. Meanwhile, as law-based governance has been extended and rights awareness has spread, more and more disputes find their way into court.

Chinese courts put people first, guided by the goal of "striving to ensure that the people feel fairness and justice prevail in every judicial case." The litigation system is designed to be accessible: filing thresholds and court fees are kept low, parties are not required to hire lawyers, and procedural timelines are strictly regulated–all aimed at delivering justice that is fair, efficient, and convenient. Yet the sustained surge in caseloads has strained this system to its limits: handling a huge volume of cases while resolving each dispute soundly and definitively has become the courts' most acute challenge.

Shenzhen is one of China's most vibrant megacities: a window on reform and opening up, and a national hub for finance, technology, manufacturing, and logistics. By the end of 2025, it had 18.25 million permanent residents, more than 22 million people under actual administration, and 4.65 million registered business entities. Economic activity is intense, and judicial demand has outgrown the supply of judicial resources. The tension between ever-rising caseloads and finite judicial resources has shifted from episodic pressure to a structural problem.

Under the traditional model, judges read case files page by page, sorted evidence and facts by hand, and drafted judicial documents word by word. Routine work consumed most of their energy, crowding out the time needed to weigh the core contested issues: judges were chronically overworked. The detection of risks such as sham litigation depended heavily on individual experience, with little systematic screening capacity. Procedural management, quality assurance, and early warning during case handling had no intelligent support, so problems were typically discovered after the fact rather than anticipated. Existing productivity tools could no longer meet the need. How to deliver fairer, faster, more convenient justice with limited judicial resources became a pressing question.



1.2 The Opportunity: A Technical Breakthrough and a Policy Framework


中华人民共和国最高人民法院
The Supreme People's Court of the People's Republic of China

《最高人民法院关于规范和加强人工智能司法应用的意见》全文（中英文版）

来源：最高人民法院　发布时间：2022-12-09 10:07:58

法发〔2022〕33号

最高人民法院

关于规范和加强人工智能司法应用的意见

为深入学习贯彻党的二十大精神，深入贯彻习近平法治思想，贯彻落实《中华人民共和国国民经济和社会发展第十四个五年规划和2035年远景目标纲要》和《新一代人工智能发展规划》，推动人工智能同司法工作深度融合，全面深化智慧法院建设，努力创造更高水平的数字正义，结合人民法院工作实际，制定本意见。

一、指导思想

1.坚持以习近平新时代中国特色社会主义思想为指导，深入贯彻习近平法治思想，坚持司法为民、公正司法工作主线，加快推进人工智能技术与审判执行、诉讼服务、司法管理和服务社会治理等工作的深度融合，规范司法人工智能技术应用，提升人工智能司法应用实效，促进审判体



> **"Promote the deep integration of technological innovations such as big data and AI with judicial work."**
>
> **– Xi Jinping**

In November 2022, LLMs achieved a breakthrough that attracted worldwide attention with their capacity for natural-language understanding, logical reasoning and text generation. Their potential in law quickly became a hot topic: LLMs can read long legal documents, summarize facts, retrieve applicable provisions and produce a structured first draft of a judicial document – precisely the routine work involved in adjudication. But general-purpose LLMs also carry serious challenges in a domain as high-stakes as the judiciary: hallucination, meaning output that looks plausible but is unreliable; the black box, meaning reasoning that cannot be explained; and safety risks.

In December 2022, the Supreme People's Court of the People's Republic of China (SPC) issued The Opinions on Regulating and Strengthening the Applications of AI in the Judicial Fields, setting the goal of building a fairly complete system for applying AI in the judiciary by 2025. The Opinions establish five basic principles: security and legality; fairness and justice; assistance in adjudication; transparency and credibility; and abiding by public order and good customs. Results produced by AI, the Opinions state, "shall be supplemental references only," and all judicial decisions are always made by judges. The document drew the framework for AI in the judiciary and gave local courts a policy basis for experimentation.

1.3 The Foundation: Paperless Case Handling, Online Litigation Services, and Knowledge Management

> **"There can be no modernization without informatization."**
>
> **– Xi Jinping**

Digitized case materials and end-to-end paperless case handling are the preconditions for applying AI in the judiciary. Shenzhen Courts' AI applications did not start from zero: they were built on more than six years of digital accumulation, moving step by step from digitization to intelligence, and from online services to intelligent assistance. Since 2017, Shenzhen Courts have progressively built the System for the Synchronous Generation of Electronic Case Files and extended paperless case handling. In 2019, Shenzhen Courts launched a mobile online litigation platform and extended services to every stage of a case. By 2023, Shenzhen Courts had established an electronic foundation covering filing, trial and enforcement, along with vast structured judicial data and mature online service channels – the practical conditions for the leap from digitization to intelligence.



1.3.1 Electronic case files and end-to-end paperless case handling

In 2017, Shenzhen Courts began building the System for the Synchronous Generation of Electronic Case Files, piloting it first at Yantian Primary People's Court, where litigation materials were scanned, catalogued, and structured automatically. In 2018, paperless case handling was extended to both levels of court across the city, and the Electronic Case-File System was connected to the mediation, service-of-process, and enforcement platforms, so that files were generated alongside the case and moved online throughout. In 2022, Shenzhen Courts' pilot on single-set electronic archiving–retaining records in electronic form only, with no paper set–passed acceptance review by the National Archives Administration of China, marking the full transition from paper to electronic files.

Six years of paperless practice left Shenzhen Courts with massive structured judicial data. Case information, evidence, hearing transcripts, and judicial documents are all preserved in electronic form, providing the indispensable raw material and the practical basis for embedding AI deeply in the work.

1.3.2 End-to-end online litigation services



In November 2019, Shenzhen Courts launched a mobile litigation service platform, the Shenzhen Mobile Mini Court, built as a WeChat Mini Program. Users register with real-name verification and, once a case is filed, the whole of the proceedings – hearing, mediation, service of process – can be conducted from a phone, for litigants and judges alike. The platform creates instant-messaging groups bringing judges and litigants together: each procedural action is visible in the group, fully logged and archivable with one click. Through this convenient entry point, litigants can check case status and progress, submit materials, attend hearings online, receive electronic service of process and take part in online mediation. Residents of Hong Kong, Macao and Taiwan, China, and foreign nationals can complete cross-border identity verification online, and cases involving large numbers of victims are supported by online registration and verification of refund accounts.

The platform has been updated continuously since launch, continuously expanding "justice for the people" digital services. In October 2025, the platform was upgraded and renamed the Shenzhen Digital Court, and now offers online proceedings end to end. Before the hearing, guided workflows help litigants submit structured data online, complete the basic case information and their claims and defenses, and present and examine evidence, so that data is captured and fixed in standard form and consolidates preparation for a concentrated, efficient hearing. At the hearing stage, litigants who choose online proceedings and

meet the conditions are served accordingly: judicial personnel log in through the court intranet, which connects securely to the internet-facing mini program used by the participants, allowing questioning and hearings to be conducted online. AI supports these hearings by analyzing the structured data gathered beforehand and putting the key facts, issues in dispute, and legal grounds before the judge. End-to-end online proceedings make hearings more focused, shorten case cycles, give litigants fuller scope to exercise their procedural rights, and make the experience of justice more tangible.

By August 31, 2026, the Shenzhen Digital Court had nearly 3 million verified registered users. It had handled more than 4.9 million cases, hosted more than 271,800 online hearings, and served more than 59.087 million documents electronically. Over 90 percent of litigation materials in the Shenzhen Courts are now exchanged and served online, cutting litigants' costs substantially, and the platform is widely accepted by litigants and by lawyers. More important still, every litigation act–submission and examination of evidence, submissions, hearings, mediation, service of process–can all be conducted online through the Digital Court, and the submitted materials exist in electronic form, laying an even firmer foundation for AI in the judiciary.

1.3.3 Authoritative knowledge bases and practical knowledge management

In recent years, the SPC has vigorously built authoritative knowledge bases–the People's Courts Case Database, Faxin (a Chinese legal knowledge platform built by the People's Court Press), and others–forming a "national legal knowledge base" that is standardized, systematic, and continuously updated. It covers everything judges need to decide cases–laws and regulations, interpretations of provisions, adjudication rules, authoritative cases, and legal views–systematically linking legislation, adjudication, case law, and doctrine, and provides a national knowledge infrastructure serving all courts in a unified manner through the dedicated court network. Thanks to this ongoing knowledge engineering led by the SPC, AI applications in the judiciary have been equipped with a national-level legal knowledge base from the very outset. This delivers four-fold guarantees for knowledge sources: authoritativeness, uniformity, completeness, and real-time updates within the dedicated network, providing irreplaceable national-grade infrastructure for subsequent training of domain-specific LLMs and agentic use of authoritative knowledge.

In parallel, Shenzhen Courts have continuously improved knowledge management at the level of judicial practice. Since 2021, they have organized senior judicial experts to codify practical expertise as adjudication guidelines covering more than 100 case types. This is not simple compilation: experience scattered across individual cases is put through a complete chain of collection, consolidation, review, conversion, and updating, and fixed as a fine-grained, highly standardized, strongly consensual body of practical knowledge that raises the quality and efficiency of case handling and makes adjudication standards more consistent. Each guideline embodies the current adjudication methods, reasoning approaches, and professional standards of judges in that field, and can be flexibly updated as practice evolves. These knowledge products – drawn from practice, feeding back into practice, and standardizing practice – guide Shenzhen judges' daily work, build specialist expertise across every field of adjudication in Shenzhen, and provide consensual corpora and knowledge tools fitted to real-world judicial work. All of these knowledge products subsequently become the core knowledge-ontology foundation of the AI applications that followed.

2. Systemic Outcomes of the Shenzhen Practice

Policy direction, a technical breakthrough and an existing digital base converged in 2023. In May 2023, at the intersection of policy guidance, technological breakthroughs, and existing digital foundations, Shenzhen Courts began planning the construction of an AI-assisted adjudication application, with strong support from the Political and Legal Affairs Commission of the CPC Shenzhen Municipal Committee and the Government Services and Data Management Bureau of Shenzhen Municipality, forming a dedicated development team of expert judges and technical staff to explore the prudent integration of LLMs and related technologies into judicial work. The SPC and the High People's Court of Guangdong Province advised throughout, keeping the work on course. In June 2024, Shenzhen Courts launched the Shenzhen Courts' AI-Assisted Adjudication System (hereinafter "the System"), which is the world's first AI assistance system spanning the entire adjudication process, achieving the pivotal leap from digital to intelligent case handling. The System has since gone through 3 major framework upgrade and close to 500 successive iterations, has incorporated agentic AI as that technology advanced, and now covers the main civil, administrative and criminal case types, assisting judges across the whole process and becoming deeply embedded in daily practice, including recent updates of autonomous AI agents.

2.1 The System's Six Core Application Scenarios

The System supports the full lifecycle of a case, covering six core modules from case-filing review to document drafting.

2.1.1 Intelligent Case-Filing Review Module





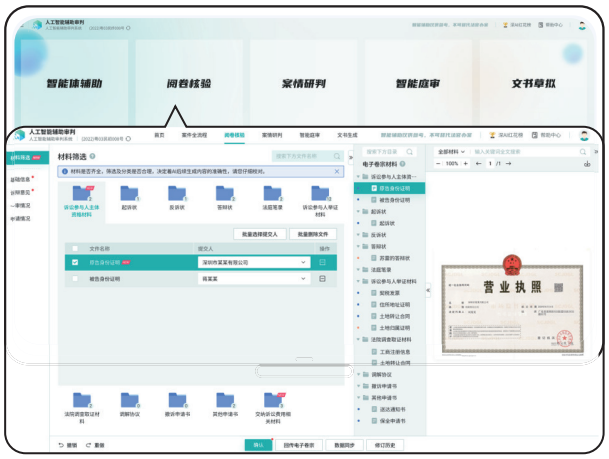
Case filing is the first gateway to judicial service. Under China's litigation system, when a party brings a suit, the court conducts a formal review of the filing materials: verifying the identity of the parties, checking that the claims are clearly stated, determining whether the case falls within the court's jurisdiction, and screening for abusive or duplicative litigation. Traditionally, this review was manual and item-by-item. The filing judge worked page by page through the complaint, the schedule of evidence, and the identity documents, comparing each against the statutory standards–slow work, and prone to inconsistency between reviewers.

Because case-filing review is comparatively structured, Shenzhen Courts chose it as the first pilot module to incorporate AI. The Intelligent Case-Filing Review Module went into service in January 2024, converting the filing standards into executable rules so that review is standardized and automated. Using intelligent document structuring, fuzzy semantic matching, and contextual inference, the module identifies, classifies, and structures the materials a party submits. Working from current law and judicial interpretations, it assists in checking all 35 filing standards, including consistency and timeliness of materials, tests the legality and compliance of the complaint, and assists in verifying jurisdiction. Cases that meet the conditions are registered in one pass; cases needing supplementation or correction of filing materials receive a single, complete notice of what is required. The module has now assisted in reviewing more than 1.3 million filings. It has made both the review of filing materials and the guidance on supplementation more precise, shortening review times and markedly reducing human oversights and divergent standards.

A natural public concern is whether AI review of case filings might reject a legitimate claim through technical error. The design principle – AI assistance with human review and final control – answers it. After the System's preliminary review, every filing decision is confirmed by a case-filing judge. Where the System marks a case as requiring supplementation, the judge conducts a second review by hand, so that no party's procedural rights turn on a decision from AI. The review standards themselves are drawn strictly from current laws and regulations and are updated whenever those change, avoiding deviation caused by outdated rules.

2.1.2 Intelligent Case-File Review and Verification Module







After the case has been assigned to a trial judge for handling, reviewing the case file is the foundation of a judge's understanding of a case, and one of the most time-consuming tasks in supporting adjudication. The volume of evidence can be very large: a sales contract dispute may involve dozens of contracts, hundreds of pages of chat records, and bank statements; the file of a construction-contract dispute can easily run to thousands of pages. Where a party is unrepresented, differences in litigation capacity often mean materials may arrive less well organized and complete, adding further to the burden. Traditionally, judges and judicial support staff read files page by page, manually excerpting key information–party names, contract sums, dates of signature, disputed clauses–into review notes. As one judge put it: "Just flipping through several hundred pages and writing down the key information takes a whole day." Worse, manual extraction invites omissions and transcription errors, and a key fact missed at this stage can distort fact-finding at the hearing and in the decision. Judges also had to switch repeatedly between systems to look up party identity information, company registration records, labour arbitration records, and other external data, further adding to the time cost.

The Intelligent Case-File Review and Verification Module turns this work from transcription into verification. It organizes and classifies litigation materials automatically, extracts the core information of the case–basic details, claims and defenses, submission and examination of evidence–and displays each extracted item against its source material on the same screen for side-by-side comparison. A judge reading an extracted item sees the original alongside it and can click through to verify it at the source. Through controlled data sharing, the module also gives judges a fuller view of external information relating to the parties and the case.

Judges no longer leaf back and forth between documents and systems to establish the facts of a case. Review becomes a matter of checking, correcting, and supplementing, and work that took a full day now takes tens of minutes. Just as important, the System scans and extracts from every document without omission, closing the gap that manual review leaves and giving fact-finding a more solid foundation.
The main technical concern here is extraction error–a contract sum read wrongly, or party A's information attributed to party B. Two mechanisms answer it: multi-source verification and same-screen comparison. The System gathers every relevant source of a given item from the case materials, filters and cross-checks them against defined rules governing judicial work, and presents to the judge first the data it assesses as most accurate. Every extracted item traces back to its exact location in the original, one click away. The System is not a black box; it is transparent and traceable throughout. Extracted information is treated as provisional data only: the judge checks, corrects, and supplements it, and final confirmation always rests with the judge. Any extraction error is caught at that point.

2.1.3 Intelligent Hearing Assistance Module





The hearing is the heart of adjudication. Its quality bears directly on the delivery of justice and, through it, on public trust in the judiciary. A civil case is usually heard once or twice, and the judge must conduct court investigation, the submission and examination of evidence, and debate within a limited time. Traditionally, pre-hearing preparation was rough-and-ready: the judge and judicial assistant read the file and set out the key points from experience, with obvious scope for omission. Mid-hearing, a party might suddenly change position or depart from an earlier account of the facts, and an unprepared judge could lose control of the pace or have to reconvene. One judge described the difficulty: "What I fear most is a party changing the factual account mid-hearing and throwing off the whole rhythm. Without preparation, it is hard to nail down all the legal facts in one sitting, and sometimes there is no choice but to schedule another hearing." Repeat hearings lengthen cases and add to litigants' costs.

The Intelligent Hearing Assistance Module covers preparation, the hearing itself, and the work that follows. Beforehand, it helps the judge identify the points to be addressed, drawing on the standard approach to that case type and on the claims and defenses in this case; checks the parties' statements for consistency and logic; sets out the facts not in dispute; proposes case-specific lines of questioning for the matters still to be established; and assists in producing a structured hearing outline. During the hearing, if something unexpected arises, the module puts the relevant legal provisions and handling guidance in front of the judge. Afterwards, it organizes the hearing transcript and analyzes it, drawing out the core issues in dispute and the state of the factual investigation.

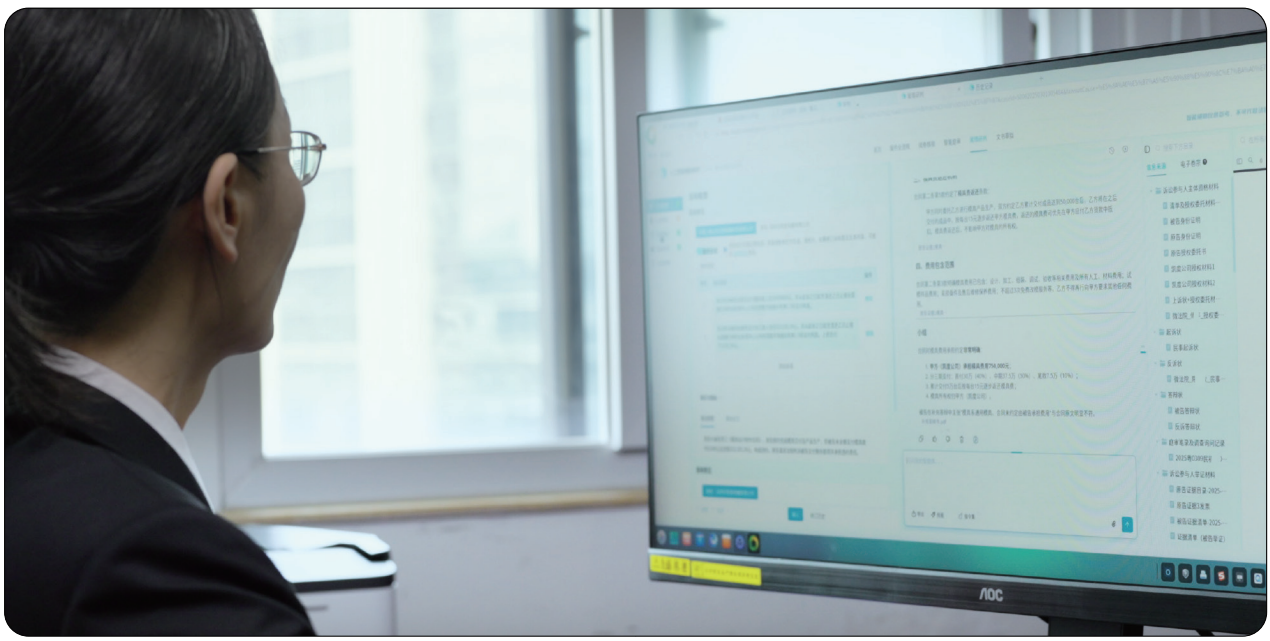
The module makes preparation more systematic and precise. A judge enters the courtroom with a full command of the contested points and is equipped for surprise evidence or contradicted testimony. In practice, judges use the outline to guide their questioning; hearings are better controlled, tighter in pace, and more sharply focused. Many cases now settle in the courtroom or reach a clear direction there, rendering hearings more substantive and purposeful, while significantly boosting the solemnity of court proceedings. Limited hearing time is used to the full, advancing the goal of resolving the dispute in a single hearing and reducing repeat sittings and procedural churn. As one judge described the module: "It is like giving every judge an assistant who works around the clock. Before the hearing opens, the System has already run a high-quality rehearsal–recommending the points at issue and flagging the key questions to be checked under each–so you walk into court with the whole case in view."

A further concern is whether AI might take over questioning at the hearing, or influence a judge's independent judgment. The design is explicit that the System is a supporting tool. What it offers is suggested questions, not instructions. A judge may adopt them, amend them, or ignore them; direction and decision at the hearing rest with the judge throughout. Legal provisions and response strategies pushed to the judge during a hearing are information only and carry no view on the outcome. The immediacy and variety of a live hearing remain, litigants' procedural rights are undiminished, and the inherent value of the hearing is not weakened by technology but more clearly displayed.

2.1.4 Intelligent Case Analysis and Assessment Module



Case analysis is the core of adjudication and the sternest test of professional judgment. From a large body of material, the judge must examine and assess the authenticity, legality, and relevance of each item of evidence, build up the facts of the case, identify the issues in dispute, analyse each in turn, and reach a conclusion. Under the traditional model, this rested almost entirely on individual expertise and experience, and it was slow and demanding. Where evidence was voluminous and the legal relationships complex, a judge would go back through the file repeatedly, noting contradictions and connections between items–work in which information is easily missed or attention misdirected. Judges also differ in knowledge and in research skills, so that in similar cases different information may inform the decision, resulting in inconsistent adjudication standards across similar cases.



The Intelligent Case Analysis and Assessment Module is the System's hub, supporting the decisive work of establishing the facts, examining the issues, drafting the reasoning, and settling the dispositive provisions. It has four components:

Summary of Claims and Defenses

The System displays the claims and defenses as reviewed and confirmed, consolidates the final statement of relief sought, and shows any amendments to the claims together with a summary of the defense.

Fact-Finding

The System assists in distilling the substance of each item of evidence, produces an evidence review that reflects the purpose for which the evidence is offered and how it was examined, extracts the factual investigation from the hearing, and – taking account of which evidence the judge has accepted – assists in consolidating the facts as found.

Key Issues for Adjudication

Working from the facts and the claims and defenses, the System identifies the issues in dispute. The judge may add, delete, or amend them, but every substantive conclusion is the judge's own. Throughout, the module provides authoritative legal knowledge, pushing the laws, regulations, reference cases, and commentary that match the issue at hand. Once the judge confirms the issues, the System assists in drafting the reasoning on each, in line with the judge's conclusions, and shows the inferential steps behind it.

Determination of Dispositive Provisions

The judge must reach a final conclusion on each claim. The System assists by generating proposed dispositive provisions from the confirmed claims and the facts of the case, and by supplying the legal grounds.

The module effectively improves the accuracy of fact-finding. By distilling the substance of every item of evidence, assessing its probative value, and forming a comprehensive review, it helps the judge examine the completeness and consistency of the evidentiary chain. In cases with voluminous evidence and complex legal relationships, it locates contradictions and gaps quickly and flags facts that might otherwise be overlooked, reducing errors of fact caused by missed information. It has also improved the reasoning in judicial decisions and the consistency of adjudication standards: whether a case is simple or complex, the same analytical standard applies, and the System prompts the judge to answer each of the parties' issues in dispute. By pushing authoritative knowledge and similar-case references at every decision point, it gives different judges the same information and the same benchmarks when handling like cases, reducing disparities arising from differences in judges' knowledge and research skills and promoting consistent adjudication standards across similar cases.

Because this module touches the core of judicial decision-making, its design had to answer the greatest concern about AI in the courts: could AI replace the judge in making decisions and reaching conclusions, and could its hallucinations about law or fact lead to a wrong judgment? The System sets three lines of defense.

First, a clear boundary of responsibility. The System supplies analysis; the judge decides. At the analysis stage, the judge may add to, delete from, or amend anything the System produces, and every substantive conclusion is the judge's own. The System offers material and analysis, and reaches no final conclusion. Judicial responsibility is therefore not blurred by the introduction of technology: the judge answers for every decision, and the System's role is confined to providing fuller information and more precise knowledge.

Second, a complete audit trail. Every analysis the System pushes and every decision point the judge passes through is recorded in full, including each confirmation, amendment, or rejection, so that the path to the decision can be retraced.

Third, the constraint of authoritative knowledge. It removes at source the risk of AI producing knowledge "hallucinations" and fabricating legal grounds. This is described under Intelligent Legal Knowledge Retrieval below.

2.1.5 Intelligent Judicial-Document Drafting Module





The judicial document is the formal record of the court's decision, and its quality bears directly on the authority and credibility of the courts. In China, judicial documents–judgments, rulings, and others–follow strict national formats and drafting requirements set by the SPC, and must contain the parties' details, the claims, the findings of fact, the reasoning, and the dispositive provisions. Traditionally, a judge wrote the whole document by hand, from formatting and data entry through to proofreading, with a great deal of time going to mechanical work. A complex judgment could take several days to draft. Data entry consumed much of that time, while the reasoning–the part that most needs a judge's attention–was compressed. Manual drafting also invites textual errors, formatting deviations, and omissions, and proofreading takes a judge's time in turn. The deeper cost is that when routine work absorbs a judge's hours, reasoning suffers; and where the reasoning is thin, the losing party in particular may feel that its case was never seriously weighed, which affects whether the parties accept the judgment and refrain from further litigation.

The Intelligent Judicial-Documents Drafting Module works from the key information the judge has confirmed at each earlier stage–the facts, the issues in dispute, the application of law, and the legal grounds–and assists in producing a complete, correctly formatted first draft in the SPC's prescribed form. The reasoning is drafted against the issues the judge has confirmed, answering each of the parties' contentions in turn. Once a draft exists, a built-in automated proofreading function detects textual errors, formatting deviations, and omissions. The System then runs a deliberately adversarial review of the

completed draft, assessing from the parties' standpoint how likely they are to accept the judgment and refrain from further litigation, which helps the judge test the reasoning again and improve the clarity of legal reasoning.

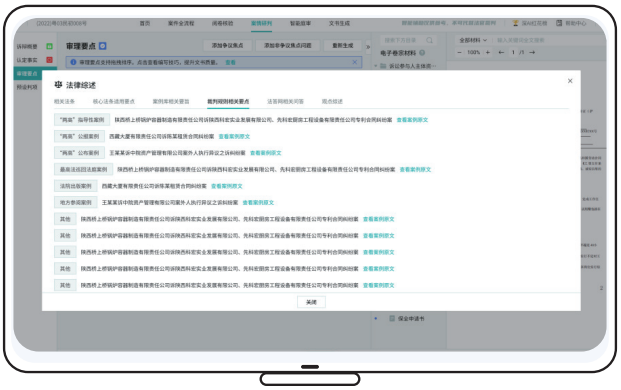
By taking over formatting, data entry, and proofreading, the System frees judges from mechanical work for the contested heart of the case. Formally, documents are produced strictly to standard, with fewer human errors and defects. Substantively, the reasoning is fuller and answers each contested point. As one lawyer reported: "My client read the reasoning in the judgment and said he felt taken seriously. As the losing party, his sense of grievance was noticeably eased." The System has now assisted in drafting 678,000 judicial documents.

At the drafting stage, the System follows one strict design rule: no one-click generation. What the System produces is a first draft; the judge is responsible for reviewing, revising, and deciding on any document drafted with AI assistance. Every finding of fact, every passage of reasoning, and every dispositive provision requires the judge's review and confirmation before the document is final. The System never produces the final document in the judge's place. Its built-in adversarial review, which anticipates from the parties' standpoint whether the judgment will be accepted, works as a form of quality control, helping the judge find weak points in the reasoning before the document is finalized.

One point deserves emphasis.

Civil, administrative, and criminal proceedings share basic features, but each has its own procedural requirements, adjudication rules, and value priorities. For ease of understanding, the five modules above are described from the standpoint of civil litigation, which accounts for more than 92 percent of Shenzhen Courts' caseload. Criminal and administrative cases are handled quite differently, and the System is designed and laid out accordingly. For criminal cases, it is built around the evidence, strengthening the analysis of the legality of evidence and of the evidentiary chain, and it provides analysis on the standard that the evidence is reliable and sufficient and the facts are proved beyond a reasonable doubt. Administrative proceedings require the court to conduct a full review of the legality of the administrative act, so in administrative cases the key issues for adjudication are not confined to the plaintiff's contentions: the System reasons across the whole act, on the principle of a complete examination.

2.1.6 Intelligent Legal Knowledge Retrieval





Legal research is indispensable to a judge's work. China's body of law is large and constantly updated, and the SPC regularly issues guiding cases, reference cases, and adjudication rules, so a judge must find the provisions, judicial interpretations, and similar-case reasoning relevant to the case in hand. Traditionally, the judge had to enter keywords and search, which raises three problems. Results depend on the precision of the keyword: if the judge's wording differs from the statutory language, important authority may be missed. Search is passive: the judge must already know what to look for, which is difficult in an unfamiliar area. And judges differ in research skills and habits, so the references obtained vary, undermining consistency in adjudication standards.

The System connects to authoritative legal knowledge bases spanning laws and regulations, commentary on provisions, adjudication rules, authoritative cases, and legal opinions, and turns them into intelligent retrieval and recommendation built for judicial use. Working with Tsinghua University, the development team introduced SAILER (Structure-Aware Pre-trained Language Model for Legal Case Retrieval) and related innovations, moving past the limits of keyword matching and textual similarity by bringing legal reasoning into retrieval itself. Drawing on the internal structure of and relationships between different bodies of legal knowledge, the System analyzes the facts, the issues in dispute, and the questions of applicable law, identifies the rules, cases, and adjudication views with a substantive legal connection to the case, and retrieves, matches, and recommends on that basis. Legal knowledge service thus moves from finding information by textual relevance to tracking the knowledge that bears on the legal questions in the case and supporting consistent standards.

Retrieval has become markedly easier, fuller, and more consistent. By pushing authoritative knowledge and similar-case references at every decision point, the System helps judges find the applicable authority quickly in complex cases, compare differences in adjudication across similar cases, and promote consistent adjudication standards. Since launch, the System has delivered 22.746 million knowledge services, making this one of its most heavily used capabilities.

The authoritative knowledge service also answers a major concern about AI in the judiciary: could AI fabricate a legal provision or a case? An LLM may hallucinate when generating text, producing plausible but non-existent provisions, case names, or adjudication views. If a judge relied on such material in reaching a decision, the damage to judicial fairness would be serious.

The constraint built into the System is strict: every legal provision and reference case must be drawn from the authoritative knowledge base, and generation by the model is prohibited. Every recommendation links to its source, so that the chain of authority, the chain of provenance, and the chain of interpretation

remain complete and traceable. A judge can see at a click which law and which article a provision comes from, and the docket number and holding of a reference case. This design substantially reduces the risk of fabricated provisions or cases by preventing the model from generating legal authorities independently., and makes AI assistance in the courts verifiable and trustworthy. In short: every legal ground the System puts forward is real, and its source can be traced.



2.2 Agents in the System





In 2026, Shenzhen Courts carefully absorbed the latest LLM-powered autonomous agent developments. On top of the white-box, rule-bound scenarios already in place, two innovations – the Human-AI Collaborative Agent and the Case Management Agent – carried assistance deeper into case-type-specific, standardized and fine-grained work, and made it more solid in the repetitive, administrative and operational work that procedure drives. Beyond the advantages agents bring in tailored assistance, efficiency and the reuse of standards, this established something more: a way of building a continuously growing "living system" in the judicial field, driven by two complementary drivers – a knowledge ontology and a technology engine.

2.2.1 The Human-AI Collaborative Agent

As described above, the System's highly customized structured interface makes clear what assistance is being given and where judicial responsibility lies, and it spares judges most of the work of writing their own prompts. But as use scaled up and judges grew familiar with working alongside AI – and came to know its strengths and weaknesses – a clear trend emerged. A competent, general-purpose assistant that scores eighty out of a hundred was no longer enough. Judges wanted a specialist partner: one that understands the context of judicial work, knows the characteristics of the case type, and interacts flexibly. Designing agents in detail around the main case types was the key to genuine human-AI collaboration.

Judges face three recurring difficulties in practice, and general-purpose AI tools are plainly limited in addressing them.

First, fact-finding assistance that does not grasp what matters in a given case type. Legal elements and adjudication rules differ enormously between case types. What matters in a private lending dispute is proof of agreement to lend and of transfer of the funds; in a labour dispute, the focus falls on establishing the employment relationship, working hours, and recognition of work-related injury. Generic tools extract elements into a uniform pattern and cannot attend to the distinctive details that decide the case, so fact-finding comes out mechanical and coarse. Unable to organize the factual elements around the legal logic and constituent elements specific to each case type, general AI produces work of uneven quality, and the judge still has to fill in much of it by hand.

Second, professional calculation – the complex legal arithmetic that general models cannot do reliably. Many case types depend on exact computation: the various forms of remuneration in labour disputes; personal injury compensation in traffic accidents; the division of marital property in family cases; interest calculated in segments at compound and penalty rates in private lending. LLMs are inherently prone to hallucination in arithmetic, so asking one to compute a figure directly invites error, and the judge must recalculate by hand.

Third, interaction that cannot accommodate a judge's individualized tasks. Judges previously had to follow the System's pre-designed steps in sequence. Those steps covered the main line of adjudication, but could not anticipate individual working methods, and human-AI collaboration remained largely a matter of combining preset AI functions with manual judicial work.

The key to solving all three, and to genuine human-AI collaboration, was detailed agent design for specific case types. Shenzhen Courts equipped these Case-Type-Specific AI agents with core skills and tools as the harness layer supporting the agents' professional capability, and added conversational interaction alongside the System's existing structured interface—upgrading both structured predefined tasks and judges' own tasks.

First, case-type-specific element extraction skills. For the major case types—private lending, labour disputes, sales contracts, theft, and others—the team built dedicated element checklists and extraction rules. Each skill embeds the structure of legal elements for that case type, locating, extracting, and presenting the key factual elements from the litigation materials. Drawing on a legal knowledge graph specific to the case type, it assembles the constituent elements, adjudication rules, and issues in dispute, and guides the drafting of the findings of fact and the reasoning—rather than applying a generic template. The Agent cuts through voluminous material with the precision of an experienced judge. What the judge sees is not generic text but professional analysis matched to the adjudication logic of that case type, and both accuracy and speed improve markedly.

Second, a self-directing Legal Calculator. For the complex calculations that recur in practice, the team developed precisely verified algorithmic tools covering overtime pay, personal injury compensation, division of marital property, and interest calculated in segments at compound and penalty rates. The work is divided: the AI Agent extracts the variables from the materials by semantic understanding, and the calculator computes. Having identified working hours, injury grade, loan dates, interest rates, and other key variables, the Agent calls the verified algorithm automatically, and the calculation is performed deterministically using a verified algorithm. Judges no longer recalculate by hand, combining AI-assisted legal analysis with reliable, rule-based computation.

Third, the "Third-Screen" Conversational Interface. The System uses a three-screen layout. On the third screen, where an agent carries out tasks a judge has set, the judge can ask questions directly and call on skills and tools; can ask for in-depth analysis of specified materials or text; can feed an agent's output back into the findings of fact, the reasoning, and other parts of the document; and can define frequently used prompts and publish the good ones to the Prompt Hub to share with all judges. This shifts human-AI collaboration from a rigid workflow of preset tasks and manual review toward a more flexible combination of predefined and judge-directed tasks.

With the addition of the Human-AI Collaborative Agent, the System strikes a good balance between standardizing case handling, implementing the judicial accountability system, and meeting judges' individual and detailed needs.

2.2.2 The Case Management Agent

Both procedural administration and substantive adjudication are essential to effective judicial work. Handling procedural work properly and promptly is not merely a matter of management: it is the foundation for protecting litigants' procedural rights and substantive justice, and an embodiment of judicial authority and seriousness. Its quality directly shapes the litigants' experience and the courts' credibility.

Case filing and acceptance, service of process and notice, scheduling of hearings, pre-hearing preparation, proofreading of documents, case closure, and archiving – each bears on litigants' procedural rights and on substantive justice. Yet these have long been the heaviest burdens judges and judicial support staff carry.

They are numerous. Measured against the statutory requirements, a single case from filing to archiving can involve more than 140 procedural scenarios requiring action by a judge or judicial assistant. A Shenzhen judge often has several hundred cases in hand at once–tens of thousands of procedural steps to track. No one can cover all of them promptly by memory alone.

They are exacting. Procedural matters nest inside one another, and each carries strict deadlines and requirements. A hearing can be scheduled only after the period for defense has lapsed; a hearing can proceed only when proof of service is complete; a case can be archived only when the transcript has been signed. An omission or delay at any point can hold up the case or create a procedural defect.

They are error-prone. With workloads as heavy as they are, judges and judicial support staff struggle to keep pace, and waiting, delay, and omission are hard to avoid. A forgotten scheduling notice may cancel a hearing; a proof of service not archived in time may block closure. These seemingly small slips, accumulating, seriously drag down the overall quality and efficiency of adjudication.

They fragment attention. A judge's time should go to establishing the facts, applying the law, and resolving disputes conclusively. In practice, a great deal of it went to tracking and executing procedural steps, squeezing out the time available for the decision itself.

In response, Shenzhen Courts made the Case Management Agent a priority in the development of judicial agents. Its principle is that the Agent follows the case and helps handle it–turning procedural work from something the judge must chase into something that comes to the judge.

End-to-end automatic awareness, so that no step is missed. The Agent senses every procedural node in a case, analysing multi-source case data to track progress in real time. The team has completed the design for five stages and 18 links of civil and commercial case handling, covering 140 operational scenarios, and has mapped more than 700 trigger conditions across nearly 500 case states. What stage a case has reached, what needs doing, what materials are still missing – the Agent identifies each and raises it, so judges no longer track every node from memory and experience.

Tasks are presented as cards, so that what is due is visible at a glance. The Agent shows outstanding tasks and progress as cards on the Case Management Dashboard, in the form judges and judicial support staff are used to. On the Dashboard, the judge sees the current progress and pending items at a glance and can issue query and operation instructions at any time. There is no more switching between systems or leafing through the file to work out the next step, because the Agent has already sorted it out.

Consolidated handling of procedural matters, from tracking to execution. The Agent does not only sense and remind; it acts. It is now connected to the underlying case-handling systems, so an instruction to the Agent can generate a procedural document, preview it, and place it in the file, or query the hearing schedule and the status of service. Case preparation, pre-hearing preparation, scheduling of hearings, review of documents, case closure, and archiving can all be handled in a consolidated, intelligent form. Once the judge confirms the decision, the Agent advances the following steps automatically along the prescribed workflow, sharply reducing manual operation.

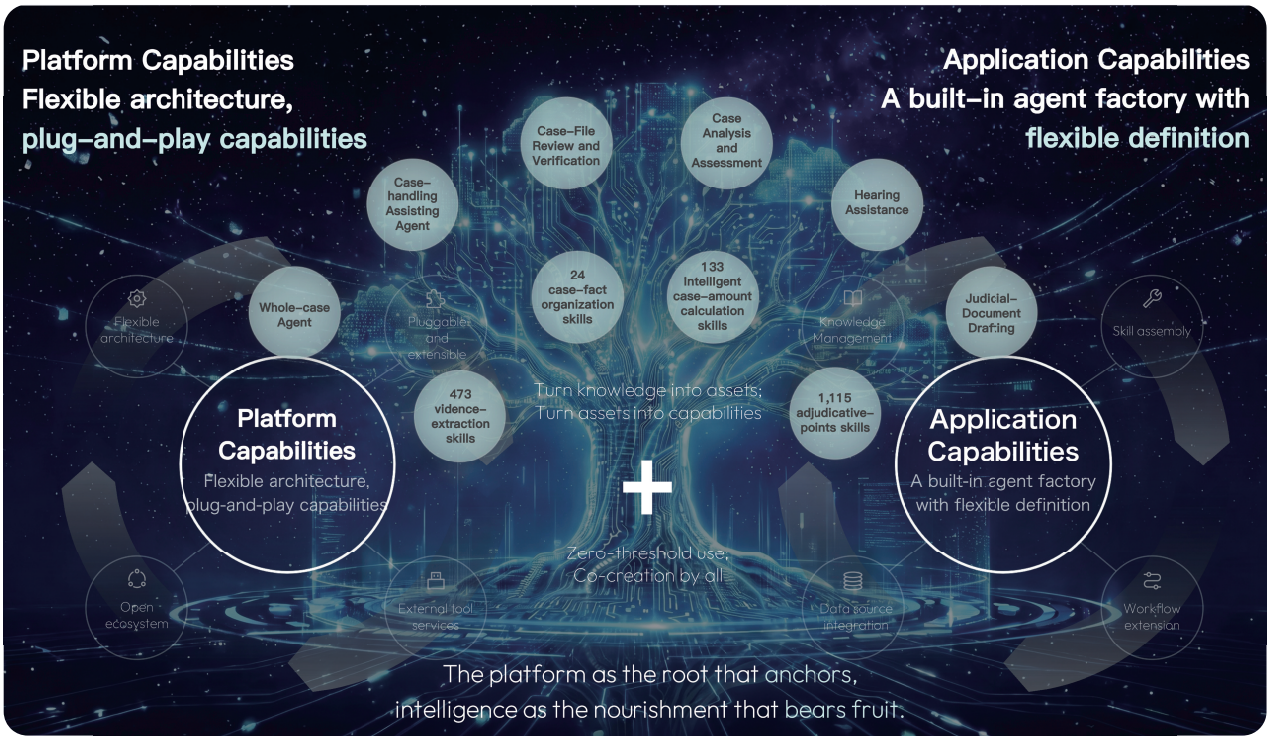
The value of the Case Management Agent to judges falls under four heads.
Relief. More than 140 procedural scenarios, nearly 500 case states, more than 700 trigger conditions–information judges and judicial support staff once had to hold and track by hand is now sensed and processed by the Agent. A judge no longer has to remember what a given case needs next; the System raises it.

Error prevention. Every procedural matter has its prescribed workflow and time limit, and the Agent advances each according to preset rules, preventing omission and delay. What must be served is not forgotten, what must be scheduled is not missed, what must be archived is not left.

Efficiency. Judges once had to check, ask, and chase. Now the Agent pushes outstanding tasks forward, the judge confirms the decision, and the System completes the rest. Procedural work moves faster, and cases move more smoothly.

Focus. This is the most important. With procedural work carried out by the Agent, judges can put more of their time into establishing the facts, applying the law, studying the rules, settling disputes, and bringing them to a definitive close–the high-value work that requires human judgment. Technology handles the routine; judges concentrate on adjudication. That was the object of building the Case Management Agent.
The Case Management Agent is the System's most recent development, with its first version released in September 2026. It continues to be refined in response to judges' feedback.

2.3 A System That Grows



AI applications in the judiciary face a further difficulty: sustaining upgrades over time. On one hand, the underlying technology moves extremely fast, and a system tied closely to one technical approach faces rebuilding from the foundations as soon as a new paradigm appears. On the other hand, professional knowledge, adjudicative logic and applicable standards in the judicial field keep evolving, and new standardized practical knowledge keeps being produced. Unless those changes are fed into the intelligent applications promptly, the applications drift away from what judicial work actually requires. At the same time, AI application in the judiciary must also be economical: Shenzhen Courts can neither afford nor seek unlimited investment in frontier technology and R&D resources. A sustainable upgrade path suited to the realities of judicial work was therefore essential.

The System accordingly departs from the traditional model of informatization, in which upgrades depend on continuous investment of resources and technology. Built for the complexity and variability of judicial practice, it is a "living system" in which knowledge can be managed, models assembled and modules exchanged – upgrading through two wheels turning together, a judicial knowledge ontology and frontier technology, so that the intelligence the courts need keeps growing.

A manageable knowledge ontology, continuously updated. So that agents keep pace with adjudication, the team built a Skills-and-Task-Chain Configuration Management Platform embodying judicial standards. The configuration platform is fully visual, low-code, and text-based, so court staff can convert the latest adjudicative logic, knowledge standards, and accumulated experience into skills and task guidance that agents can use. Through chain configuration, task definition, and Harness Engineering, this knowledge is delivered precisely to the AI agents' capabilities, turning the continuous management of judicial knowledge into a knowledge-ontology tool that responds instantly and directs precisely, available to judges and agents alike, and enabling adjudication standards and best practices to be applied consistently at scale. Besides drawing real-time updates of legal knowledge from the national authoritative knowledge bases, Shenzhen Courts have built on the knowledge ontology already formed through earlier knowledge management—the adjudication guidelines and related products—and organized expert judges to make implicit adjudicative standards explicit and transferable; knowledge that passes professional review is then converted technically. More than 1,000 certified items, including skills embodying adjudicative standards, procedural workflows, and proprietary adjudicative algorithms, have been integrated so far. Feeding this body of best practice into the System continuously keeps it current and growing.

A modular and configurable technology engine, continuously updated. So that judicial work keeps sharing in the advance of frontier AI, and on a common base of shared data and knowledge assets, the System uses loose coupling, modular deployment, and evolvable microservices. Every engine in the multi-model engine integrating general-purpose and domain-specific LLMs can be swapped as required, and every application module deployed flexibly.

Large models can be replaced as needed. Around the specific judicial scenarios and the Harness Engineering design, Shenzhen Courts organized nearly 200 judges from different fields to write more than 5,000 high-quality questions in a judicial test bank, producing a distinctive evaluation framework for AI capability in legal practice, together with a cycle of tracking, testing, replacement, release, and feedback. Any secure and suitable LLM or agent can therefore be substituted into a specific task or service—a layered pattern of continuous improvement in Harness Engineering, continuous updating of the knowledge ontology, and continuous iteration of the intelligence engine.

Modules can be decoupled and improved separately. Element extraction, case-file review, similar-case recommendation, hearing assistance, document drafting, and the other sub-modules are loosely coupled and separately deployed, so each can be improved rapidly on the basis of judges' feedback, addressing specific problems step by step while the System as a whole stays stable.

New services can be plugged in easily. The System keeps an open, modular, microservice-based architecture, so any technical solution or mature application that serves judges better can be connected quickly. The System currently integrates frontier technologies from six industry-academia-research partners, genuinely drawing on the best available talent and sustaining an R&D ecosystem aimed at the best performance outcome, the best development cost, and the best user experience.

3. Innovations of the Shenzhen Practice

The demonstrative significance of Shenzhen Courts' AI-Assisted Adjudication System lies not in any single technical breakthrough, but in turning AI from a laboratory concept into a productivity tool spanning the entire adjudication process–a complete, globally leading scheme for AI application in the judiciary.

3.1 The First Systematic, End-to-End Judicial AI Application

Worldwide, experimentation with AI in the judiciary is not rare, but the great majority of practices remain at a single task or limited scenarios. Against this backdrop, Shenzhen's System is the world's first systematic AI application covering all litigation types and running through the entire adjudication process.

3.1.1 Full coverage of common proceedings and case types

Version 1.0, released in June 2024, already covered the main civil and commercial case types. Coverage then expanded: the administrative module in September 2024, the criminal module in December 2024. The System now covers more than 90 percent of case types across civil, administrative, and criminal proceedings, so a judge has AI assistance in the great majority of cases. This is the world's first extension of judicial AI from a single stage to every type of proceeding and every case type, and no judicial AI application anywhere has matched that breadth.

The significance lies less in the numbers than in how closely the System is fitted to the legal logic of each type of case. Civil cases turn on the validity of contracts and liability for breach; criminal cases on the completeness of the evidentiary chain and consistency in sentencing; administrative cases on review of the legality of administrative acts. The three differ markedly in approach, evidence rules, and document structure. Shenzhen Courts brought in judges from almost every field to build the data and write the skills, converting the adjudication logic of each case type into executable intelligent skills. AI assistance is therefore not generic text processing but precise support matched to the professional requirements of each case type.

3.1.2 Full coverage of substantive and procedural work

The System's other breakthrough is that it covers procedural and substantive work together. Handling procedural matters properly and promptly is the baseline guarantee of litigants' procedural rights and of substantive justice. Substantive work–fact-finding, application of law, reasoning–determines the outcome. Within the System, AI helps judicial support staff organize files, verify information, handle procedural matters, and check and format documents, and helps judges with the core judgments of fact-finding, application of law, and reasoning.

The System is built for general applicability and embedded in the courts' existing systems, connecting closely to the Electronic Case-File System and the Shenzhen Digital Court platform. The Digital Court handles online submission and standardized capture of litigation materials at the front end; the Electronic Case-File System stores and archives materials and handles procedural matters; the System performs intelligent analysis and decision support on that basis. Together, the three form a single chain–capture at the front, processing in the middle, intelligence at the back–with no fragmentation between systems and no data silos.

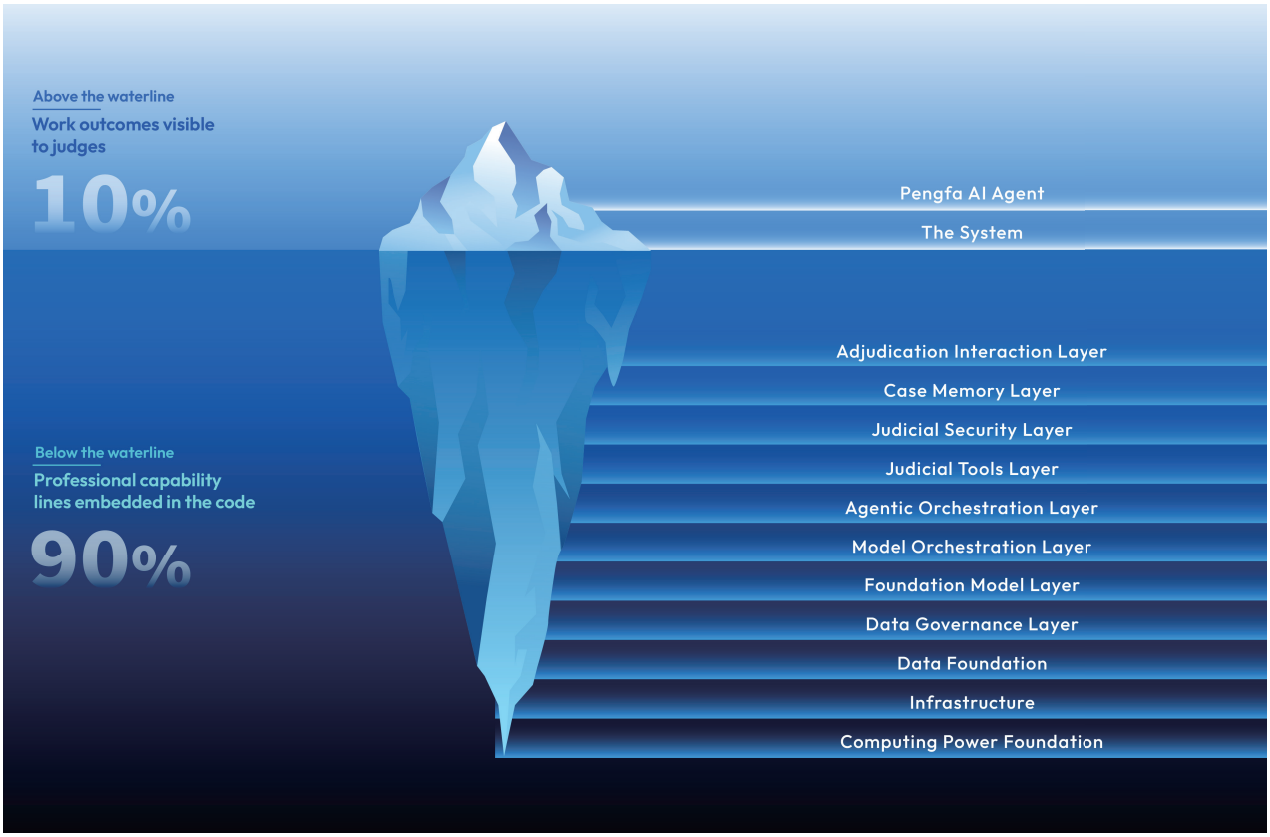
The System runs through the five main stages of case-filing review, case-file review and verification, hearing assistance, case analysis, and document drafting. Alongside them, the Process-Driven Agent senses every procedural node automatically, analyzes multi-source case data and, following the judge's decisions and the prescribed workflow, handles case preparation, pre-hearing preparation, scheduling of hearings, review of documents, case closure, and archiving in a consolidated, intelligent form.

Such a complete coverage makes AI assistance not an occasional tool but infrastructure integrated with every stage of adjudication. From filing to closure, the System accompanies the judge and carries out more of the routine work, so that a judge's attention can go to the core questions in dispute and to the exercise of judgment. This is what AI applications in the judiciary have long aimed at and not yet achieved anywhere: not replacing judges with AI, but using AI to release judges' professional capability–technology handling the routine, judges concentrating on adjudication.

3.2 The First Domain-Specific Harness Engineering for the Judiciary

LLMs reason and generate powerfully, but a model by itself does not equal productivity. A model that converses fluently, yet does not know what stage the case is at, which tools to invoke, or which decisions must be returned to the judge for confirmation, cannot truly enter the adjudicative workflow. The industry calls the engineering approach that solves this problem "harness engineering." Its core idea is "humans steer, agents execute": the model thinks, reasons and acts, while the harness controls it through to a delivered result. Building constraint mechanisms, feedback loops, workflow control and continuous improvement around an agent turns unpredictable, error-prone AI into a stable, reliable production tool.

For the judiciary this matters more than elsewhere. Judicial work demands accuracy, explainability and clear allocation of responsibility far beyond ordinary industry requirements. An answer that is roughly right is unacceptable in court; it must be right. Shenzhen Courts, therefore, did not simply transplant a general-purpose harness. Judicial experts led the overall architecture, starting from how judicial work is actually done: mapping the full process and defining where AI could add value, where rules needed restructuring and where processes needed rebuilding. The result is a Judiciary-Specific Harness Engineering Framework embodying judicial standards. It carries three core tasks.



3.2.1 Turning models into judicial productive capacity: from talking to working

An LLM can reason, write, and summarize, but on its own, it does not know how to use tools safely, how to remember the state of a case, how to observe the rules on adjudicative authority, or how to recover from error. The model thinks; the harness controls how it acts. In a judicial setting, that division means the model supplies legal understanding and reasoning, and the harness organizes that capability into deliverable work products that conform to the adjudication process.

By building the multi-model engine integrating general-purpose and domain-specific LLMs with authoritative knowledge base services, the framework assembles context dynamically according to the specific task in hand. The System handles four key steps in a predefined workflow: organizing case memory, admitting legal knowledge, calling certified skills, and incorporating the judge's decisions. It does not simply pose a question to a model and wait for an answer; it behaves like a trained expert assistant that knows what stage the case has reached, which skills to call, which authoritative legal knowledge applies, and which steps must return to the judge for confirmation.

Specifically, Shenzhen Courts broke the adjudication process into five stages–case-filing review, case-file review and verification, hearing assistance, case analysis, and document drafting–and equipped each with its own skills and knowledge services. When an agent executes a task, the harness orchestrates the order of tool calls, manages case memory (isolated by case and cleared by task, so that information from one case cannot contaminate another) and enforces execution boundaries, so that every output conforms to adjudication logic and adjudicative standards. This conversion from model capability to judicial productive capacity is the value of Harness Engineering: not making the model cleverer, but making it work more reliably in a judicial work environment.

3.2.2 Preventing hallucination: layered defenses from input to output

Hallucination is inherent to LLMs: an LLM may produce content that looks plausible and is not reliable. In a judicial setting, that means inventing provisions, fabricating case facts, or offering the wrong view on the applicable law. Were such content to enter a judicial document, the consequences would be grave. The judicial domain-specific Harness Engineering therefore sets four layers of defense from input to output.

First layer: knowledge admission–only authoritative knowledge enters the reasoning. The System strictly prohibits AI from generating legal provisions or case texts. All knowledge on which legal reasoning rests must come directly from certified authoritative knowledge bases, and false legal information in litigation materials is identified and flagged to guard against knowledge contamination. All LLMs are restricted from recalling provisions from their own parameters ; they retrieve and cite from a governed authoritative knowledge base–closing the principal channel through which hallucination enters.

Second layer: input and output are filtered by layered guardrails. Multiple guardrails, built from rule bases and discriminative models, identify, intercept, and handle content on the way in and the way out. Direct user input is filtered for safety and screened for jailbreak attempts, and obviously manipulative instructions will be refused. Third-party information an agent obtains by calling external skills and tools is filtered for injected instructions and isolated by format.

Third layer: judge-in-the-loop decision-making chain with human confirmation throughout. At key nodes, the System sets up repeated verification, escalating review, knowledge admission, and human decision-making . Multiple LLM-powered agents assist across the process and enforce quality checks at important points, but every substantive conclusion must be made expressly by the judge–the acceptance and evaluation of evidence, the findings of fact, the application of law, and the decision itself. The System is strictly prohibited from offering substitute decisions in the form of a recommended judgment or a one-click generation of conclusion, and any automated chain that would bypass human judgment is ruled out.
Fourth layer: visual traceability throughout-white-box, trustworthy assistance. The System keeps the source of case information, the original text of the knowledge, the judge's decisions, and the generated output visible and traceable at every point. Every analytical conclusion can be traced to the evidence, the legal provision, and the knowledge base text it rests on, and the judge can examine, confirm, or reject each in turn. This transparency across the whole chain turns the System from black-box output into white-box, trustworthy assistance. The judge sees not only the conclusion but how it was reached.

Together, these four layers contain the hallucination risk of LLM applications and make AI in the judicial setting usable, useful, and easy to use.

3.2.3 Achieving cost-effective compute: engineering optimization instead of stacked hardware

Inference and post-training consume a great deal of compute, and courts as public institutions have limited resources. Simply scaling up the size of models would be unaffordable and would not necessarily deliver the professional precision and stability judicial work requires. Shenzhen Court’s Harness Engineering proves another route: rather than pursuing maximum model performance, optimize the domain-specific environment in which the model runs.

Precise judicial chains of thought and context engineering to reduce unnecessary inference. By assembling context dynamically, the System injects only the case information and legal knowledge relevant to the task at hand, keeping unrelated material out of the reasoning. This improves accuracy and reduces the compute each call consumes.

Modular service called on demand. The functional modules–element extraction, case-file review, similar-case recommendation, hearing outlines, document drafting, and the rest–are deployed independently and invoked only at the corresponding stage, rather than running one large unified model throughout. Calling only what is needed makes the most of limited compute.

Tiered compute through the multi-model engine integrates general-purpose and domain-specific models. The System integrates different types of models with different scales of parameters and selects the appropriate tier automatically according to task complexity: smaller language models handle simple routine work, and larger models are called only for complex legal reasoning assistance. This tiering cuts total compute consumption substantially without compromising output quality.

Authoritative knowledge bases in place of model memory. The System does not rely on an LLM remembering every provision and case; it retrieves them in real time. Parameter counts, therefore, need not expand indefinitely to hold ever more legal knowledge, and compute can be directed at reasoning rather than storage.

This approach–engineering optimization instead of stacked hardware–allows the System to cover more than 90 percent of cases across the three types of proceedings on limited computing power. It supports an important conclusion: in a vertical domain, the refinement of Harness Engineering does more to determine real-world results than parameter count. It is also a replicable experience in applying AI deeply where computing resources are limited.

任务分解后消耗的算力及成本
Computing Power and Cost Consumption for Same Tasks



3.3 The First Judicial AI Security Framework Deployed in Production

As digital courts and judicial technology continue to advance, the in-depth integration of artificial intelligence brings unprecedented capabilities alongside inherent risks. Without a rigorous, systematic security governance framework, AI may fail to bolster judicial fairness and could instead become a hidden threat to public trust in the judiciary. Shenzhen Courts have adhered to security-oriented design from the system's initial development stage and continuously reinforced relevant safeguards thereafter, fostering a secure, trustworthy, and stable operational environment. Three core measures underpin this robust security system.

3.3.1 Closed-loop data on a dedicated network

LLM-powered AI has raised the quality and efficiency of case handling sharply, and with that comes a severe test of data protection and privacy. Court files and judicial documents contain a great deal of personal identity information, commercial secrets, and socially sensitive data. If transmitted or used for model training via public clouds or the Internet, such data is exposed to unauthorized capture, malicious interception, and covert extraction. The risk touches not only litigants' personal rights but the foundation of judicial credibility. Under the guidance of the SPC , as well as the planning and management of Shenzhen municipal authorities, Shenzhen Courts deployed the System entirely on the dedicated network for political and legal affairs, which is physically isolated from the internet. This approach effectively eliminates underlying security risks and protects the private information involved.

Top-level planning and institutional standards. Faced with emerging security challenges, the SPC and Shenzhen municipal authorities have issued rigorous management standards and technical specifications from judicial domain and local governance perspectives. Building upon this foundation, Shenzhen Courts adopted unified review and deployment to ensure all domain-specific models are developed, trained, and applied within a robust regulatory framework, precluding loss of technical control and improper usage.

3.3.2 Security by design

Adjudication is a highly rigorous scenario, and flaws at any technical layer, including ambiguous top-layer functional logic, middle-layer black-box inference algorithms, and vulnerable bottom-layer models or AI agents , may cause erroneous decisions and undermine judicial credibility. Isolated patches cannot address the complexity of judicial practice, so security must be embedded throughout the entire development lifecycle. The system adopts full-stack security-by-design principles, covering its top-level architecture and knowledge framework, middle-layer trustworthy algorithms and reasoning chains, as well as the underlying orchestration platform, foundation models, and intelligent agents. This layered protection shifts security governance from passive remediation to proactive built-in assurance. Having passed stringent national assessments, the system has obtained the highest-level agent security certification issued by the China Academy of Information and Communications Technology under the Ministry of Industry and Information Technology, setting a technical benchmark for secure and trustworthy AI in judicial adjudication.

Upper layer: algorithms and context chains. The System uses explainable AI and a strictly tamper-proof context chain, guarding against prompt injection along the whole chain, so that every step from evidence extraction through logical inference to document generation is traceable to its source. This approach addresses logical drift and context contamination in the atomic use of LLMs: every task in the chain is examined for the provenance of its inputs, minimizing the risk of hallucination and malicious injection from the root, as well as improving the explainability and traceability of the reasoning.

Base layer: orchestration platform, foundation models, and agents. The system eliminates vulnerabilities arising from fragile infrastructure and agent over-authorization. As the underlying infrastructure underpins overall system stability, flaws in foundation models, privilege escalation on the orchestration platform, and out-of-control agent behavior during multi-step reasoning may all pose systemic risks. To address such hazards, Shenzhen Courts has built a fully self-controllable model orchestration platform, adopted security-hardened foundation models, and invoked Zero-Trust Architecture and dynamic privilege isolation into agent interaction workflows. This ensures that autonomous agent operations, including retrieval, reasoning, and summarization, are consistently constrained within strict security boundaries and sandboxed environments.

3.3.3 Security by specialized mechanism

As the System moved from development and testing through pilot operation to large-scale deployment, its critical question shifted from whether the AI worked to whether operation was secure and stable. To build a technical management model suited to judicial use, Shenzhen Courts established specialized security mechanisms covering development, deployment, release, and operations.

At the development level, DevSecOps is applied throughout the stage, embedding code security audits, sensitive-data masking standards, and static vulnerability scanning at every node, so that security is written into the code from the outset and the ground for intrusion and latent defects is removed at the source.

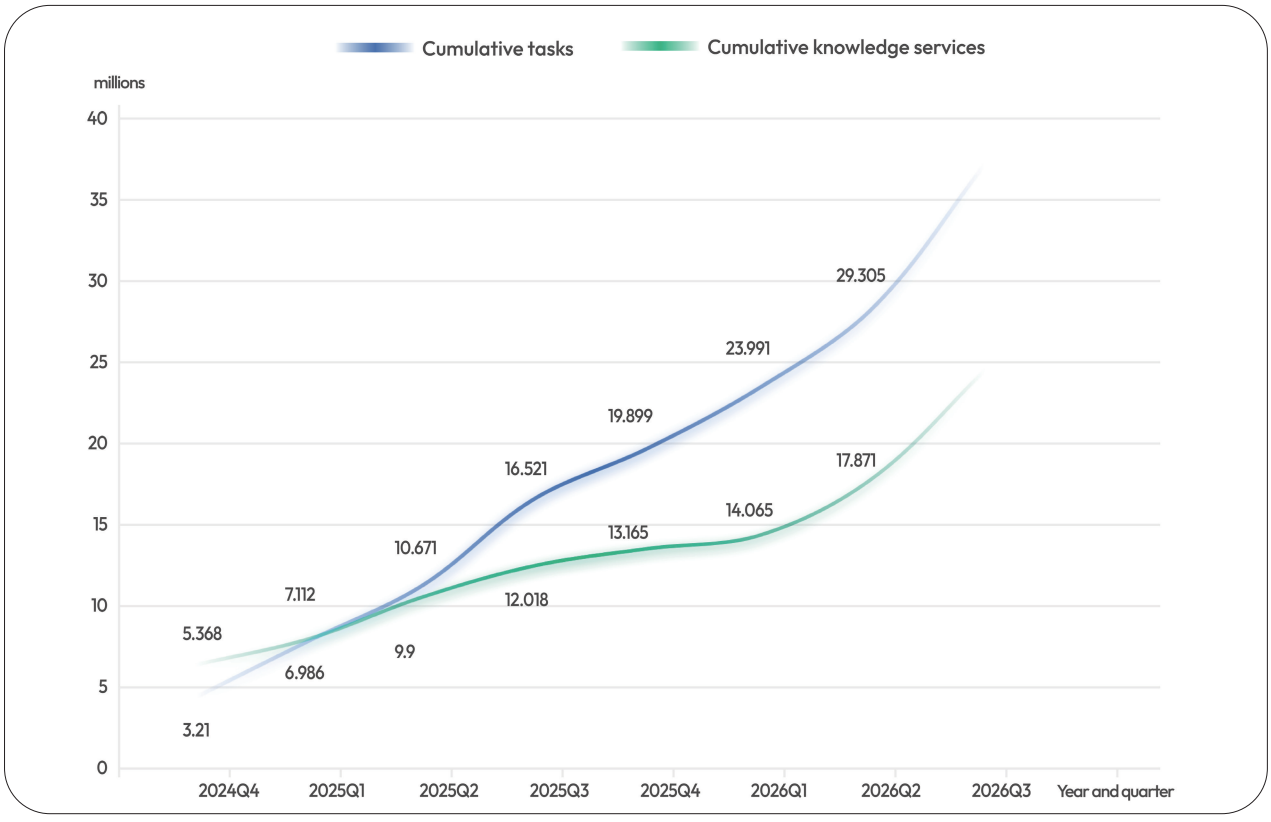
At deployment and release level, strict release-approval and user-acceptance-testing are applied. Any new function or version must pass stress testing, security assessment, and judicial domain-specific testing in an isolated user acceptance environment. Only after completing judicial and technical reviews and obtaining official approval can any update be released to production, ensuring that every iteration is stable, controlled, and reversible.

At the operational structure level, the independent and isolated architecture eliminates traditional security risks caused by scattered and overextended operational privileges. Supported by least-privilege enforcement and dedicated channel isolation, alongside round-the-clock automated log auditing and abnormal traffic monitoring, the system achieves real-time operational visibility, second-level response to minor anomalies, and precise fault localization.

At the security mechanism level, full-time security officers are vested with independent review authority and veto power. They oversee the full lifecycle from system development to operational deployment, and conduct regular security assessments and emergency drills. Acting as the full-life-cycle security sentinel of the system, security officers bridge management gaps in human-AI collaboration and ensure full institutional compliance with all security requirements.



4. Practical Application Outcomes of the Shenzhen Practice



609476 case-file reviews

268678 hearings

678003 judicial documents

Since its large-scale launch, AI assistance has become a stable and effective new quality productive force in the judiciary across both levels of court in Shenzhen. From launch to 31 August 2026, LLMs and AI agents had completed 35.451 million tasks. The System assisted with 609,476 case-file reviews, 268,678 hearings and 678,003 judicial documents, supporting more than 700,000 cases in depth and easing the imbalance between judicial workload and staffing. Upholding quality as the paramount focus, the System delivered 22.746 million legal knowledge services during the adjudication process, enabling improvements in both the quality and efficiency of adjudication.

Behind these figures lie technology and institutional reform working in concert. Shenzhen Courts have aligned their workflow with the System's operation, refined adjudication management and specialized adjudication mechanisms, and fully mobilized judicial personnel. Reform and the System moved in step, and the results are measurable. In terms of quality, the appeal rate fell by 48 percent and the rate of reversal or remand for retrial fell by 65 percent. In terms of efficiency, Shenzhen Courts concluded 773,828 cases of all kinds in 2025, an average of 744 per judge – 249 more than the previous year; the average number of cases concluded per judge per month was 60 percent higher than before the System went online.

Judges' own accounts give these changes concrete form. A judge at the Pingdi Dispatched People's Tribunal of Longgang Primary People's Court wrote: "The System helps judges return to judging... A judge is no longer troubled by the anxiety of whether the work can be finished, and can concentrate instead on how to do it well." A judge in the No. 2 Civil Division of Nanshan Primary People's Court observed: "For me, the System is like the navigation you use when driving; it has become an indispensable assistant in my work... However rough the road of adjudication, with technology behind us, we can travel faster and more steadily towards fairness and justice." Voices from the front line show that the value of AI-assisted adjudication lies not only in taking over routine work but in freeing judges' attention for the core adjudicative jobs, balancing efficiency against quality.

The "AI + adjudication" practice of Shenzhen Courts has won the SPC's full recognition and strong support, and has become a template for judicial intelligent transformation in China.

The strength of the technology developed in-house, and the results it has produced in practice, have brought Shenzhen Courts wide recognition across China: more than 200 domestic institutions, courts at every level among them, have come to Shenzhen on dedicated study visits. Visitors have consistently found that the framework applies systems-engineering principles throughout, fits the realities of the work, and transfers readily to other judicial and government service settings. In April 2025, the Digital Shenzhen Joint Innovation Center was established, with a branch center set up within Shenzhen Courts in parallel. In September 2025, the SPC established the National Court Judicial AI Application Joint Innovation (Shenzhen) Base, which is the only base in China.The Base built a high-fidelity, full-scenario hardware and software experimental environment in Shenzhen to pursue the vision of "innovation in one location, shared nationwide". Leveraging these two major innovation platforms and their dedicated hardware and software experimental environments, Shenzhen Courts undertakes the R&D, application, feedback, and iteration of intelligent innovation services for the judiciary.

Under the innovation base, new functions have been developed–drafting judicial recommendations, similar-case search reports, in-depth legal research, and rapid drafting documents for simple cases–and fully incorporated into the Unified Case-Handling and Office Management Platform for People's Courts Nationwide. These are now in trial operation at courts across the country and will be rolled out nationwide. At the same time, under the SPC's guidance, Shenzhen Courts are consolidating a standards framework for judicial AI application and building the theoretical framework for its use, settling widely accepted principles in the form of technical application guidelines and judicial ethics guidance, so that what Shenzhen has piloted becomes practical standards for others.

The System's launch and deep application have drawn keen attention at home and abroad. The innovation not only leads and demonstrates within the Guangdong-Hong Kong-Macao Greater Bay Area; with its originality and advancement, it has won wide international attention and recognition, gradually becoming a showcase card of China's achievements in building the rule of law. The Journal of Legal Analysis – an open-access journal published by Oxford University Press on behalf of Harvard Law School – published: "How do judges use large language models? Evidence from Shenzhen" hails the Shenzhen practice as a pioneering example: "the first court in the world to routinely and extensively use generative AI" in its work. Delegations from the legislative, executive, and judicial organs of Hong Kong and Macao have made repeated study visits; more than 20 delegations of judicial leaders from 20 countries–including Brazil, Vietnam, Singapore, and Portugal–have visited on more than 20 occasions. The delegation from the Supreme Court of Singapore commented on the practice as "where the world is moving forward."

5. Experience from the Shenzhen Practice

Building and operating the System was a landmark undertaking with no chart to steer by. Nowhere in the world was there a mature precedent for the systemic application of AI across a court's work in the judiciary. Shenzhen was entering genuinely uncharted territory, with no paradigm to borrow, no technical route to reuse, and no evaluation standards to consult. Every step forward probed the unknown; every decision weighed judicial prudence. Through that exploration, understanding advanced in a rising spiral, through three successive shifts: from belief in a single "almighty" model, to a systematic view of engineered assembly; from treating model parameters as the source of knowledge , to establishing white-box reasoning with authoritative knowledge bases as the sole source; from the "standardized model" orchestrated by experts, to the "open ecosystem" defined by judges themselves. Three major trial-and-error lessons, three shifts– each not a simple negation, but a deeper re-understanding of the relationship between AI and adjudication.

The first shift:from domain-specific model primacy to engineering-led delivery.

In the early days of development in 2023, Shenzhen Courts tested the large-model bases available on the market; none met the demands of production deployment, so they turned to building their own vertical legal-domain large model, feeding statutes, judicial interpretations, judgments, Q&A corpora, and the like into large-scale post-training and instruction fine-tuning to give the model a thorough command of legal knowledge. The domain-specific model far surpassed general models in provision retrieval, similar-case matching, and legal-terminology understanding. But a deeper problem then surfaced: it remained essentially a Q&A tool, unable end-to-end to meet the practical needs of adjudicative-assistance tasks of every kind. Practice has proved that, however finely trained, a single model alone cannot carry a case from the facts to a finished judicial document on its own. This recognition drove a fundamental route adjustment: Shenzhen Courts moved beyond reliance on a single, all-purpose model and adopted an engineering approach–workflow orchestration, expert-built skills, and engineered assembly. The model

was no longer the direct producer of answers but one component in a larger system: judicial experts defined tasks, distilled tools, and designed explainable chains; engineering linked the components and enforced model orchestration and safety constraints; and judges reviewed and confirmed at key points, keeping judge-in-the-loop decision-making. This was the first conceptual leap, from a model-centered to an engineering-centered approach.

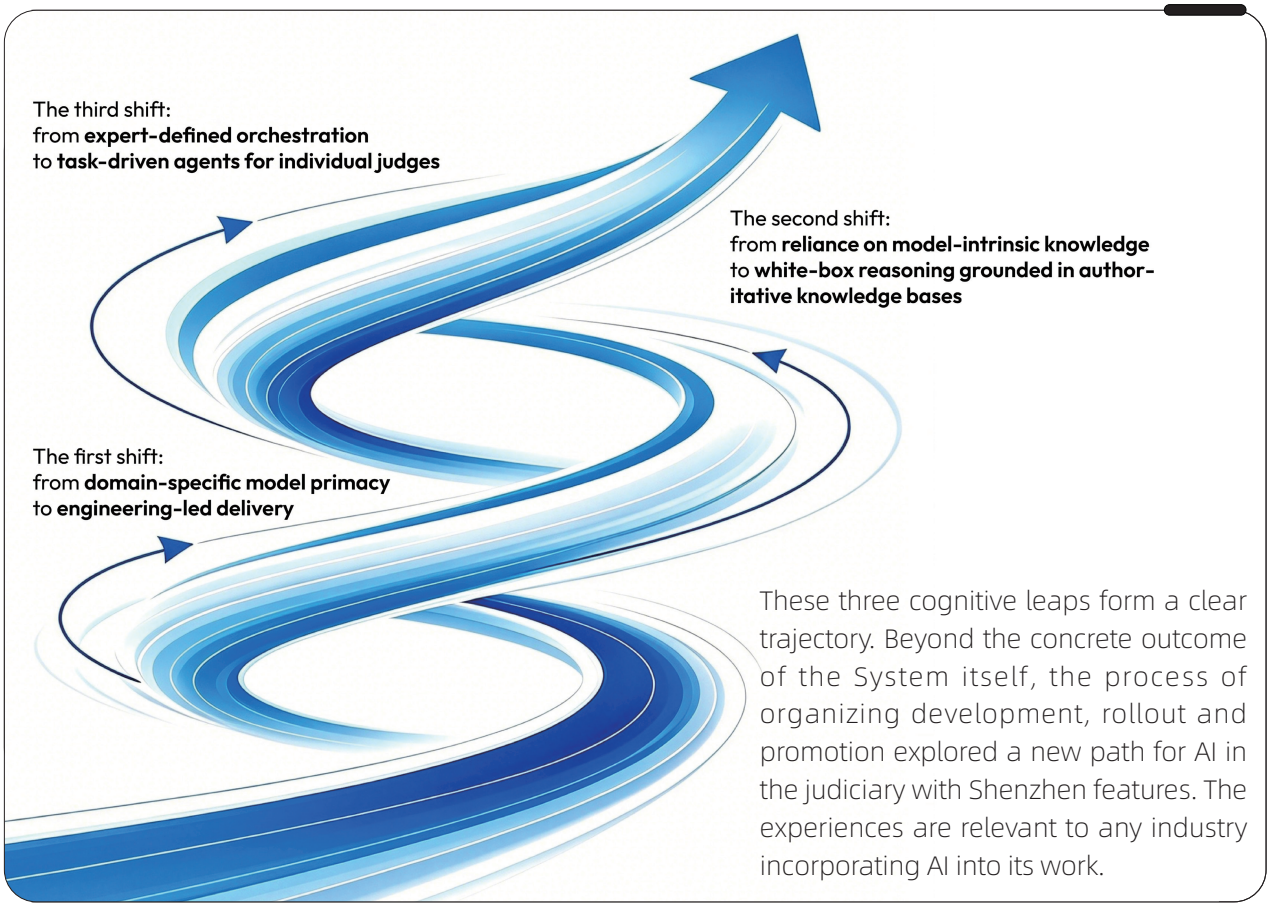
The second shift: from reliance on model-intrinsic knowledge to white-box reasoning grounded in authoritative knowledge bases.

Once the domain-specific LLM had shown stronger legal capability than general-purpose LLMs, the team briefly assumed the obvious next step: since the model had internalized vast legal knowledge, why not let it power all knowledge tasks directly, embedding safety mechanisms and explainability modules inside the model? Several approaches were tried—chain-of-thought constraints, verification against legal rule graphs, and adversarial training to catch hallucinations, each achieving phased results in small-scale experiments. In production, however, one fact kept reasserting itself: as long as an LLM draws knowledge from its own parameters, hallucination cannot be eliminated. For the high-stakes judicial scenario, this is intolerable. An LLM may produce non-existent statutes, outdated judicial interpretations, or fabricated case information, and it typically does so with a high confidence that bears no relation to accuracy. If false information entered a judicial document, it would shake the judicial foundation of deciding cases "according to facts and law." The team therefore drew a red line: LLM parameters should never serve as a direct source of legal knowledge. The System's security foundation was rebuilt around it. The SPC's continuously updated knowledge bases became the sole authoritative source of knowledge, and every algorithm module and agent was restricted to bridging—between the knowledge base and the case, and between the knowledge base and the judge's task. The domain-specific legal model may understand, retrieve, compare, organize, and carry knowledge, but never supply it. A white-box knowledge admission mechanism requires that any legal knowledge entering a judicial document or a reasoning step be confirmed by the judge and verified against the knowledge base, making the reasoning chain white-box and explainable. Every step is traceable to its source: what the judge sees is not a black-box guess about the LLM, but a verifiable path linked directly to the authoritative knowledge base's original text. This three-part architecture – authoritative knowledge base, knowledge admission, white-box reasoning – remains the System's anti-hallucination foundation today.

The third shift: from expert-defined orchestration to task-driven agents for individual judges.

With the engineering approach settled, Shenzhen Courts assembled senior judges and engineers to standardize each agentic workflow around identified best practices, aiming to give judges a ready-made operational framework. Early feedback after launch was good, but deeper use exposed a dilemma of standardization versus personalization. No single expert-defined logic fitted every case-handling scenario, and individual needs surfaced from every direction: some judges preferred a particular similar-case retrieval logic, some had a distinctive way of organizing evidence, and some had a characteristic style and argumentative preference in the reasoning of their decisions. However refined, a uniform process or workflow designed from the vantage point of experts could not cover them all. The deeper realization was that over-reliance on expert orchestration might contradict the very goal of bringing intelligence into judicial work.

A judge's creativity, discretionary reasoning and individual exploration are not incidental to adjudication; they are part of what adjudication is, and they drive both the growth of legal knowledge and substantive justice in individual cases. An AI-assisted system that suppresses them would betray the meaning of assistance. That recognition produced the most important redesign of all: taking the capability of judicial domain-specific Harness out of experts' hands and providing it directly to individual judges. The System now has incorporated professionally certified skills at scale, fitting these atomic units of best practice into each judge's own planning chains and tasks.There is also a Prompt Hub, letting judges share and reuse one another's personalized working methods: a judge can encode personal experience as prompts, graft them onto the large model's capabilities to create a bespoke supporting tool, and other judges can browse, borrow and adapt those prompts to suit a different style. The result is a shift from one uniform chain for all to personalized workflows for individual judges: within a single secure, controllable engineering framework, every judge can configure and invoke AI assistance to match a personal approach and habits. Personalization is no longer an edge case; it is the System's core design assumption. The new design holds a strict balance between open technology and judicial red lines: one-click generation of final judicial documents remains strictly prohibited, and the expert-defined standard chains must still be followed step by step – but how each intelligent task is carried out is now for the judge to define.



These three cognitive leaps form a clear trajectory. Beyond the concrete outcome of the System itself, the process of organizing development, rollout and promotion explored a new path for AI in the judiciary with Shenzhen features. The experiences are relevant to any industry incorporating AI into its work.

5.1 Concentrating Resources: Higher-Court Guidance and Local Support

No single court or tech company can independently develop AI applications that satisfy every judicial requirement. Shenzhen Courts encountered four limiting conditions: the judicial sector lacked a dedicated LLM; the court system had no internal computing infrastructure; legal knowledge resources were spread across separate institutions; and technology evolved far more rapidly than the update cycles for standard IT projects. To resolve these issues, Shenzhen Courts operated under professional guidance from the SPC and the High People's Court of Guangdong Province, relied on leadership from the CPC Shenzhen Municipal Committee and support from the Municipal Government, and built an open platform via collaboration and resource sharing. It brought together top-tier technology, high-fidelity legal knowledge, and maximum available computing power to serve judicial work.

5.1.1 Local support: coordinated development

Shenzhen is a city built on innovation, a national standard-bearer in the high-tech industry, and it offers institutional backing, technical capacity, and an innovation ecosystem that few places can match. The municipal authorities attached great importance to the AI-assisted adjudication project from the start of development, treating AI in the judiciary as a priority breakthrough in building Shenzhen into a pioneering "AI+" city. The Political and Legal Affairs Commission of the CPC Shenzhen Municipal Committee and the Government Services and Data Management Bureau of Shenzhen Municipality guaranteed the funding, technology, data, and channels the work required. More than 100 million yuan of public funding and up to 230 PFLOPS of computing power were committed over time, truly embodying the principle of "to pool resources to accomplish major tasks" and thereby ensuring a sustainable foundation for the System's development. The Government Services and Data Management Bureau provided end-to-end oversight for the whole project. During the feasibility-study stage, it acted on the special-economic-zone ethos of enterprise, pioneering action and diligent work. It convened expert reviews, released formal opinions as quickly as practicable, assigned computing resources, and assisted teams to resolve operational hurdles. At every project milestone, the bureau delivered technical guidance and strategic direction. It allocated resources and rolled out a full set of policy supports via the Digital Shenzhen Joint Innovation Center and municipal government programmes to sustain ongoing system upgrades and security safeguards.

5.1.2 Data: a corpus approaching a trillion characters

When courts deploy AI, "knowledge hallucination" represents the top risk for judicial bodies in every jurisdiction. To address this risk at its root, developers must ground System's operations on authoritative, reliable legal knowledge bases. With the SPC's support, Shenzhen Courts cooperated with the People's Court Press to integrate the People's Courts Case Database, the Legal Answers Website (Fadawang), and the Faxin Platform, giving the System a specialist legal corpus approaching a trillion Chinese characters together with real-time linkage to the SPC knowledge database, and raising the precision of legal research and rule application to the maximum extent.

5.1.3 Technology: a full stack of Chinese-developed hardware and software

AI in the judiciary requires substantial computing power and technical support, while the security and self-reliance requirements of judicial data mean the System cannot depend on external compute or technical services. The System was developed and integrated by Shenzhen DIB Technology Co., Ltd., which holds national confidentiality qualifications to develop privacy-sensitive systems; with joint efforts from all participating parties, the top-level judicial application architecture and knowledge framework, middle-layer trustworthy algorithms and context pipelines, and underlying model orchestration platform, foundation models, and agents are all independently developed. In cooperation with Huawei Technologies Co., Ltd., Shenzhen Courts took the lead in adapting the entire System to Huawei's Ascend AI computing chips, becoming Huawei's global computing model site for AI-powered judicial applications. The System has been granted four national invention patents and one design patent across its human–AI interaction layer, the application design layer, the agentic orchestration layer, and the Domain-specific Harness layer–securing both technical independence and security, and the stability and sustainability of operation.

5.2 Rules First: Setting the Direction and Holding the Line

When Shenzhen Courts first explored the uncharted territory of deploying LLMs for judicial workflows, it avoided rushing directly to write code. Instead, it first researched applicable governance rules. It consolidated technical principles, judicial ethics and usage-security requirements into a set of foundational rules. Those rules guided the development work and drew the security red lines that could not be crossed. Putting rules first kept the work on course and avoided detours. In 2023, Shenzhen Courts set out four development principles–position AI as a supporting tool; integrate AI with the judicial work; implement judicial accountability in the loop; ensure data security throughout the process–which guided the first attempt anywhere to integrate LLMs across the whole of adjudication.

In 2026, focusing on the technical foundations and potential risks of LLM-powered autonomous agents, Shenzhen Courts joined the Shenzhen Research Institute of Big Data (SRIBD), the Institute for Internet Judiciary, Tsinghua University (THUIJ), and the Faculty of Law, The University of Hong Kong to publish a four-point negative list–no delegation of decision-making powers; no input of ungoverned data; no use of uncertified skills; no bypassing of security guardrails–together with the world's first Guide for Using AI Agents in the Judicial Domain, now in version 2.0, which added practical guidance for four stages of agent use from the judicial user's standpoint: define the task, supervise the process, verify the output, and harness the loop.

Talking about the way out is worth less than walking it out. The framework developed by Shenzhen Courts through practice both secures the sound use of AI in a domain as rigorous as justice and contributes general rules and standards for AI in high-stakes domains. The SPC strongly endorses this framework. Its wider relevance is reflected in the attention it has attracted beyond mainland China: delegations from 20 countries across Asia, Europe, the Americas, and Africa, as well as from Hong Kong and Macao, China, all recognized these rules for their practical value following their field visits.

5.3 Adjudication-Led, Cross-Disciplinary Development

Integrating AI into judicial work is not as simple as transporting off-the-shelf technology into court facilities. Shenzhen Courts confronted a core friction point: technical teams lacked deep familiarity with adjudication logic, while judicial staff lacked comprehension of technical concepts and technical limits. Traditional development workflows assign technical developers to lead projects while judicial stakeholders only participate as end users. This pattern consistently yields systems that function technically but prove cumbersome for real courtroom work. To address this challenge, the system has adopted an innovative methodological development paradigm. This paradigm features four defining attributes: problem-driven design, adjudication-led development, an integrated and flat organizational structure, and agile iteration.

5.3.1 Problem-driven: starting from real bottlenecks

The first step was not writing code but sitting judges down to map out what they actually did each day and which tasks consumed the most time. The adjudication team broke the five stages of adjudication, including case-filing review, case-file review and verification, hearing assistance, case analysis, and document drafting, into 85 workflow steps.For each step, the team flagged pain points and bottlenecks: which information required repeated verification, which text segments needed manual re-entry, and which evidence-sorting tasks consumed substantial time yet delivered limited value. This workflow inventory defined the task categories that AI agents could perform.

Judges and judicial support staff headed up the core R&D team. This team set the System's end-to-end workflow and technical architecture, and it crafted the detailed workflows for AI-assisted adjudication functions. The team compiled detailed documentation of adjudication rules, translating tacit judicial experience into explicit, executable knowledge and rules. The System is therefore designed by judges, tested by judges, and built to serve judges–rendering it explainable, testable, verifiable, and trustworthy in operation, and bridging the gap between technically viable and genuinely useful AI. By addressing challenges arising from their own case practice, judges developed a strong sense of ownership, fostering a problem-driven innovation culture, with continuous efforts to tackle bottleneck issues.

5.3.2 Adjudication-led: adjudication logic drives system design

In the whole development process, Shenzhen Courts held to the principle that judicial practice leads. AI serves adjudication, not the reverse.

Judicial professionals made all final decisions for core functional requirements, process, redesigned roadmaps and priority rankings. The technical team converts judicial requirements into viable engineering deliverables, and technical teams hold no authority to rewrite established judicial workflows. For role allocation, judicial professionals analysed relevant legal provisions and adjudication rules, setting out procedural standards and working methods, while the technical team translated those requirements into executable system logic and adopted knowledge-base validation and other technical approaches. Together, they developed practical workflows aligned with real adjudication practices. The collaboration entailed more than a simple hand-off of requirement documents. Both teams accepted accountability to ensure every function matched real-world adjudication realities. In terms of organization, Shenzhen Courts mobilized 40 judges to participate in development, over 300 judges in testing, and more than 3,000 judges and judicial support staff for usage and feedback. This ensured the System evolved in line with judicial work logic rather than being driven by technical guesswork.

5.3.3 Integrated and flat: Breaking down professional silos for integrated R&D

Conventional IT projects are typically split along departmental boundaries and organized by hierarchy: one party defines requirements, another delivers implementation, and leadership provides sign-off. The longer the chain, the lower the efficiency. Shenzhen Courts broke this model by forming an integrated

team of full-time and part-time members. The team cuts across departmental and hierarchical lines and adopts flat development. Each group assumes a clear share of responsibilities to reduce conflict, duplication, and idle waiting.

The core team brought together judicial professionals and technical talent to work side by side. Shenzhen Courts entrusted young practitioners with substantive responsibilities: Young judicial professionals in Phase One core team averaged just 32.6 years of age. The court appointed these staff to fill judicial, technical and "product manager" roles. They coordinate design cycles, testing work and feedback flows across project subgroups and function like an entrepreneurial working group.Weekly meetings set targets, resolved bottlenecks, and adjusted implementation paths. Placing decision-making authority with end users of the tools shortened the decision chain and boosted execution efficiency. Judges and judicial support staff working full-time and part-time on development were expected to achieve a two-way leap: they needed to master not only legal expertise but also the basic logic of AI, including the first principles of LLMs as well as the limits and application potential of such technology in judicial scenarios. In turn, the technical team received training on legal concepts and judicial procedures to gain insight into frontline casework. This deep cross-disciplinary integration fostered a genuine joint development community.

5.3.4 Agile Iteration: Continuous Evolution Driven by a Practice-Based Flywheel

Waterfall development pattern carries structural disadvantages for AI-era projects, for judicial and for other domains alike. Teams freeze requirements too early, run long delivery cycles, involve users too little, and defer most risk identification until final stage project phases. These shortcomings prevent waterfall workflows from keeping pace with rapidly evolving judicial policies and changing judicial needs. Shenzhen Courts moved beyond one-off planning and prolonged delivery cycles in favour of an iterative approach of "build, use, and concurrently refine".

The team did not aim to deliver a completely polished finished product on the first try. Instead, it ran short, fast iteration loops and deployed new capabilities directly into live judicial workflows. From official development kickoff to production go-live, the project spanned seven months, and teams built 1.0 version of the System from the ground up. After rollout, Shenzhen Courts set up a permanent feedback channel to collect and address specific difficulties raised by judges on a daily basis, releasing a new version every week based on frontline data and feedback. In the first month alone, 877 issues were identified and resolved, and daily problem reports dropped quickly from 183 to 69. In just over two years, the System went through three major architectural releases and over 500 agile micro-iterations and refinements. From Version 1.0 in June 2024 to Version 2.0 in December 2024 and Version 5.0 in August 2026, its coverage expanded from key civil and commercial case types to administrative and criminal modules, extending into core tasks including evidence analysis and fact-finding. It now covers more than 90 percent of case types.

Throughout the development process, practical effectiveness served as the sole evaluation criterion. Project teams measured two core outcomes: whether judges found the System useful, usable and intuitive; and whether workflows delivered measurable gains in case-handling quality and efficiency. Evaluation relied not on technical indicators, but on judges' real-world operational experience and tangible improvements in judicial practice.

Consistent delivery of practical, value-adding outcomes fostered judges' trust in the System and confidence for human-AI collaborative workflows. In this way, the System has evolved into a dynamic living platform that supports continuous iteration and incremental capability growth.

5.4 Growing the Ecosystem: Changing Mindsets and Deepening Application

System launch marks merely the starting point, rather than the end. The real test lies in whether judges are willing and competent to use the System, and capable of driving its continuous improvement through in-practice application. Even a technically sophisticated system becomes a waste of advanced technology if deployed yet left idle or operated without effective feedback mechanisms. Shenzhen Courts focused on building a sound ecosystem for judicial AI to take root and serve as a regular, long-term component of judicial work, advancing efforts across four key dimensions: mindset transformation, capacity building, institutional management reform, and talent cultivation.

5.4.1 Changing mindsets: from being told to use it to wanting to use it

When AI enters adjudication, judges most often respond with scepticism. Can it deliver reliable output? Who is responsible if something goes wrong? Until such concerns are answered, even highly-capable systems cannot achieve broad adoption.

Shenzhen Courts adopted a peer-led communication strategy: let actual users share unvarnished real-world experience. Judges who participated in development piloted the System first, then travelled to adjudication divisions and primary courts to share their experience, outlining concrete use-cases where the tool delivered value, areas where performance fell short, and best practices for maximum operating efficiency. At the same time, the court's intranet launched a special-purpose AI application column named "Pengfa AI". The column has published 103 curated firsthand judge accounts and examples of AI-assisted casework. Inside the System itself, an honour-roll feature delivers public recognition for active users.

The measures relied on peer influence grounded in real case experience rather than top-down mandates. Once judges found that the System genuinely cut down repetitive routine labour and freed more time for their core "judging and deciding" judicial responsibilities, their mindsets shifted organically.

5.4.2 Building skills: making sure everyone can use it well

After judicial staff accepted the value proposition, the next priority became building practical operating skill. Shenzhen Courts delivered large-scale, multi-level hands-on training programmes for all judges, judicial support staff and clerks.

A "Qiming" (Enlightenment) Instructor Team was drawn from development team members and skilled users, teaching in person at the adjudication divisions and the primary courts.These instructors were not technical specialists. They were serving judges who deployed the System for their own live casework, and they taught practical operational know-how rather than abstract theory. In 2025, Shenzhen Courts partnered with the Shenzhen Federation of Trade Unions and the Human Resources and Social Security Bureau of Shenzhen Municipality to organize a citywide competition in AI-assisted adjudication. More than 4,100 judges and other court personnel took part, marking the largest participation occupational-skills competition ever held within Shenzhen. The competition identified a talent pool of adjudication professionals with strong AI proficiency. Organizers widely circulated the winning practitioners' working methods, lifting baseline skill levels across the whole staff. Competitive participation motivated self-directed learning; improved learning in turn promoted real-world tool adoption. This shift transformed capability-building away from one-way passive training toward interactive peer-to-peer knowledge exchange.

5.4.3 Reforming management: giving human-AI collaboration institutional support

New technology cannot unlock its complete potential if organizations leave legacy management frameworks unchanged. Shenzhen Courts reformed management on two fronts.

Specialized adjudication of categorized case types. In recent years, courts have been diverting simple and complex cases into different processing tracks. Fast-track judges process batches of simple matters and issue streamlined judicial documents. This model has greatly lifted operational efficiency and mitigated pressures stemming from heavy caseloads and constrained judicial staffing.

Once the System is deployed at scale, Shenzhen Courts can re-examine case-diversion workflows through the lens of judicial specialization.Teams can allocate matters that formerly moved through fast-track procedures to receive careful full-fledged adjudication. Judges can include complete legal reasoning inside every judicial document. The workflow preserves processing efficiency while delivering higher consistency across adjudication standards. Shenzhen Courts therefore revised its case-diversion model. It groups defined case categories together and fully restores specialized adjudication arrangements. In June 2025, Shenzhen Courts issued the Action Plan for Further Strengthening Specialization and Improving Adjudication Quality and Efficiency. Under this plan, assigned dedicated judges handle categorized case-types that make-up 60 percent of each division's incoming caseload. This establishes a specialized operating mechanism: deep on field familiarity delivers both speed and quality. The reform improves ruling quality, shortens proceeding timelines and strengthens overall case outcomes. Real-world practice validates that specialized adjudication of categorized case types both secures efficiency and promotes greater consistency in adjudication standards.

A virtuous-circle model of adjudication management. Shenzhen Courts defined performance benchmarks to create a virtuous circle in case handling operations with a balanced target: resolve roughly the same volume of cases that newly arrive, without sacrificing adjudication quality. The Courts dynamically administer this benchmark system with built-in early warning triggers and corrective exit mechanisms. Primary courts, adjudication and enforcement divisions, and judicial teams are given substantial operational autonomy: units that meet the standard set their own tailored plans, while units that do not are supported by the adjudication management department together with the personnel and inspection departments. This reform shifts management mindsets in three dimensions: teams move from reactive firefighting toward proactive governance; they shift focus away from purely final-outcome oversight toward full-process monitoring; and oversight expands beyond standalone adjudication-management functions to implement joint supervision together with personnel and inspection teams. These changes strengthen motivation across the entire adjudication workforce.

5.4.4 Accumulating talent and collective intelligence: turning best practice into system capability

The aim of cultivating this ecosystem is not to produce a small cohort of AI enthusiasts. Its core goal is to harvest insights from the full breadth of judicial talent, transforming institutional wisdom into the System,so that collective intelligence accumulates and spreads in the organization.

The System adopts the emerging "agents + skill tools" technical paradigm. Seasoned judicial staff map frequent case categories, construct professional skill tools, finalize document templates, design production workflows, and refine supporting mechanisms. It distills judicial best-practices and formal standards into highly-focused standardized intelligent tool modules.Once teams develop an efficient skill purported by one individual judge, every user across the organization can access that skill. Individual experience thereby converts into shared collective institutional capability.

To sum up, four pillars work in tandem: shift mindsets so judges want to adopt the System; deliver training so judges possess required operational skill; reform management systems to embed formal institutional backing for human-AI collaboration; and accumulate talent so teams continuously consolidate best practices. Collectively these elements build a complete ecosystem where judicial personnel participate across the full value chain: development, live usage, feedback and iterative refinement.This institutional foundation will not be lost because of technical upgrades or staff turnover. The ecosystem supports self-renewal and long-term self-sustenance, and this feature delivers the fundamental guarantee needed for regular, sustained judicial AI development.

6. A Vision for Digital-Intelligence-Native Justice

6.1 Long-Term Challenges and Contemporary Opportunities in Global Judicial Work

Today's world is undergoing a historic process in which AI is profoundly reshaping modes of production, ways of life, and forms of social governance. New-generation AI technologies–LLMs, AI agents, embodied intelligence–are advancing rapidly, with applications spreading across government services, financial regulation, healthcare, education and research, and social governance. AI is no longer merely a standalone tool or auxiliary software; it is becoming infrastructure that helps organizations run, re-engineer their processes, and build new capabilities. As an important component of the national governance system, the judicial field must inevitably face and proactively respond to this technological transformation.

From a global perspective, the pressure of rising caseloads against limited judicial resources is not unique to China. Rather, it is a structural challenge for modern court systems in many countries. For example, in recent years, official reports issued by the United States federal courts, the Indian court system, Brazil's National Council of Justice, and the Organisation for Economic Co-operation and Development have all shown that case backlogs, court congestion, continuously increasing workloads for judges, and lengthening case-processing cycles have become major governance issues confronting modern judicial systems. How to continuously improve the supply capacity of justice under limited judicial resources has become a key question jointly explored by courts around the world.

The United States federal courts have established a "judicial emergency" system specifically to measure the severity of pressure on judicial resources. When judicial vacancies, combined with caseloads, exceed certain standards, a court is deemed to have entered a "judicial emergency," which serves as an important basis for prioritizing the supplementation of judicial resources at the federal level. Among these standards, a district court may be designated as experiencing a judicial emergency where the weighted caseload per judge exceeds 600 cases, among other circumstances. As of July 2026, the official website of the U.S. courts listed 13 courts nationwide as being in a judicial emergency. In some of these courts, the annual weighted caseload per judge had exceeded 800 cases, reflecting the long-term high caseload pressure borne by certain courts.

In OECD Economic Surveys: Brazil, the Organisation for Economic Co-operation and Development summarized the long-standing judicial pressure faced by Brazilian courts as high litigation demand, case congestion, slow case processing, and heavy judicial workloads. The report noted that courts in different Brazilian states generally face a situation in which large numbers of pending and newly filed cases are concentrated among limited judicial resources, and that court organizational capacity has become an important factor affecting case backlogs. According to official nationwide court case statistics released by the National Council of Justice (CNJ), Brazil's national judicial body, approximately 10 civil and commercial cases were newly filed per 100 residents in 2018. The number of newly filed civil and commercial cases nationwide increased from approximately 5 million in 1990 to 29 million in 2017, representing an almost sixfold increase. Affected by persistently high litigation demand, the average disposition period for civil and commercial cases in Brazilian courts of first instance is approximately 850 days. In response, the OECD recommended measures such as optimizing court organizational structures, strengthening the allocation of judicial resources, improving alternative dispute resolution mechanisms, and reducing unnecessary litigation demand, so as to ease court congestion and improve the operational efficiency of the judicial system.

In India, the National Judicial Data Grid, established by the e-Committee of the Supreme Court of India, publishes real-time information on the operation of cases in courts across the country. As of March 13, 2026, district courts and subordinate courts nationwide had a total of 49,956,834 pending cases. The data grid also uses case backlog monitoring as an important decision-making tool to support judicial management and resource allocation, providing data support for courts at all levels to continuously monitor case operations.

These developments show that huge caseloads, insufficient supply of judicial resources, and continuously increasing adjudication pressure are practical challenges commonly faced by modern judicial systems. Case backlogs not only increase the workload of judges, judicial support staff and other judicial personnel, affecting their physical and mental health as well as the sustainability of their careers, but may also lead to longer case-processing cycles, fluctuations in adjudication quality, and increased litigation costs for parties. These consequences may, in turn, weaken judicial credibility and the people's sense of gain to fairness and justice. In a sense, justice delayed is justice denied. How to continuously improve the supply capacity of justice, safeguard adjudication quality, and expand the coverage of judicial services under the condition of limited judicial resources has become an important issue that must be answered in the modernization of judicial work.

Against this background, using AI to assist judicial work and improve the supply capacity and operational efficiency of justice is the inevitable result of the interaction between technological development and judicial demand, as well as a historical trend. AI is capable of rapidly processing massive amounts of

information, identifying and extracting key facts, assisting with similar-case retrieval and legal research, conducting rule matching and knowledge reasoning, and generating reasoning for judgments and other structured texts. It can take over repetitive, routine tasks, freeing judges and staff to focus on work that demands judgment, enabling judges and judicial support staff to devote more energy to core tasks that more directly involve judicial judgment.

At the same time, AI is not merely a technical tool for improving work efficiency. To a certain extent, it can also serve as an organizer, provider, and auxiliary analyst of judicial knowledge. By systematically integrating legal norms, adjudication rules, experience from similar cases, and judicial policies, it can provide judicial personnel with more comprehensive, timely, and precise knowledge support and decision-making references. AI-assisted justice is not simply about "using machines to replace people." Rather, it is about using technology to allocate resources better, strengthen adjudication support, and improve litigation services.

Meanwhile, proactively and properly introducing AI into judicial work is also a practical imperative for protecting the integrity, fairness, and security of the judicial process. Against the backdrop of the rapid popularization of AI, members of the public, lawyers, corporate legal staff, and even judicial personnel may all come into contact with and use various general-purpose AI tools. If courts do not proactively build internal AI-assisted systems that conform to the internal rules of judicial work, meet security requirements, and possess regulatory capabilities, judicial personnel may still, in actual work, use external general-purpose models to assist in searching for materials, organizing case facts, or drafting documents.

Such ad hoc, scattered, and unregulated use may give rise to multiple risks. First, there are data security risks: case materials, personal information, trade secrets, state secrets, and confidential information relating to adjudication work may be leaked during uploading, retrieval, or storage. Second, there are risks concerning the authenticity of content: general-purpose models may generate fabricated legal provisions, incorrect cases, distorted facts, or inaccurate reasoning. Third, there are risks concerning adjudicative responsibility: where the extent of AI involvement is not transparent, the traceability of the judgment formation process and the attribution of responsibility may be affected. Fourth, there are risks of value bias: the training data, algorithmic mechanisms, and output preferences of external models may not be fully consistent with the intrinsic values of the judicial system and legal system. Fifth, there are risks arising from regulatory absence: courts may find it difficult to promptly discover, assess, and correct the potential harms caused by improper use of AI.

Therefore, rather than passively facing the external impact of AI , it is better to proactively incorporate AI into the court governance system and adjudication work system. Through institutional design, technical architecture, access control, log retention, risk assessment, and accountability mechanisms, courts can realize the standardized, secure, and controllable application of AI in the judicial field.
History repeatedly shows that breakthroughs arrive fast, but institutions take a generation or more to absorb them. Time waits for no one. Faced with breakthroughs in AI, the judicial system must proactively accelerate this process. On the basis of following the internal rules of judicial work and the operating rules of adjudicative power, and on the premise that judges prudently exercise adjudicative power and bear judicial responsibility, it is necessary to actively promote the deep integration of AI and judicial work. This means fully realizing the positive value of AI in improving efficiency, optimizing services, and assisting decision-making, while prudently managing the security risks, ethical risks, and procedural risks it may bring. This is the inevitable choice for courts to keep pace with the development of the AI era and respond to the public's demand for justice.

6.2 Creating a "Digital-Intelligence-Native" Mode of Judicial Work

Shenzhen Courts believe that, in the process of integrating AI with judicial work, the auxiliary role of AI must always be upheld. AI may assist judges, but it cannot replace judges; it may provide reference opinions, but it cannot make adjudicative conclusions; it may participate in information processing, but it cannot substitute for judicial judgment; it may improve case-handling efficiency, but it cannot weaken procedural justice; it may promote consistency in adjudicative standards, but it cannot rigidify discretion in individual cases. The design, deployment, and use of any AI-assisted system must be based on the fundamental premise that "judges make independent judgments in accordance with the law and bear responsibility in accordance with the law." At every stage of case handling, judges' final decision-making authority, duty of review, and responsibility-bearing mechanisms should be clearly defined and strengthened, so as to ensure that AI always remains in an auxiliary position that is explainable, traceable, reviewable, and controllable.

On this basis, judicial work methods should be transformed with a more forward-looking concept: taking "digital-intelligence-native" as the goal, and rethinking and reconstructing the operating model of modern justice. "Digital-intelligence-native," meaning the organic integration of data-native and AI-native approaches, refers to a model in which, in the design of court business systems, technical systems, work processes, and management mechanisms, AI is no longer treated as a single tool added after the fact. Instead, data governance and AI capabilities are embedded from the outset as core elements into the overall architecture, operational processes, and institutional norms of judicial work. Its essence is not to add an "intelligent button" outside traditional working methods, but to carry out process reengineering, capability reshaping, and rule optimization for adjudication, enforcement, litigation services, adjudication management, judicial transparency, workforce development, and other areas of work, based on the capability characteristics of AI and the internal rules of judicial work.

Creating a "digital-intelligence-native" mode of judicial work means that court work requires comprehensive digitization, structured data, and standardized processes. High-quality data is the foundation of AI-assisted justice. Without comprehensive, accurate, standardized, and usable data, AI cannot effectively play its role. Courts should continue to promote online case handling, synchronous generation of electronic case files, intelligent court hearing records, digitized enforcement information, and online litigation services, so that case materials, procedural nodes, adjudicative outcomes, enforcement progress, adjudication management information, and other data can form unified, standardized, and governable data resources. At the same time, courts should strengthen the development of judicial data standards and improve full-lifecycle management mechanisms for data collection, cleaning, labeling, storage, retrieval, sharing, and deletion, thereby providing a reliable foundation for knowledge injection, model training, effectiveness evaluation, and continuous iteration of AI-assisted adjudication systems.

Creating a "digital-intelligence-native" mode of judicial work means re-examining the work content, job responsibilities, and collaboration models of judicial personnel with AI capabilities at the core. Under the traditional adjudication model, judges and judicial support personnel need to spend a great deal of time reading case files, extracting information, conducting searches, making comparisons, summarizing, drafting, and proofreading. As AI capabilities improve, a considerable portion of these tasks can be initially processed by AI systems, while judicial personnel will assume more responsibilities in instruction design, result review, judgment confirmation, quality control, and responsibility bearing. For example, in the production of legal documents, judges will shift from word-by-word drafting to designing the reasoning, reviewing the argumentation, and making the final call on content. In evidence review, AI can help extract evidentiary elements, mark points of dispute, and identify contradictions or corroborative relationships

among evidence, but the admission of evidence and determination of facts must still be made by judges in accordance with the law. In the application of law, AI can assist with legal provision searches, similar-case recommendations, summarization of adjudicative rules, and prompts regarding sentencing ranges, but legal interpretation, rule selection, and case-specific balancing should still be completed by judges.

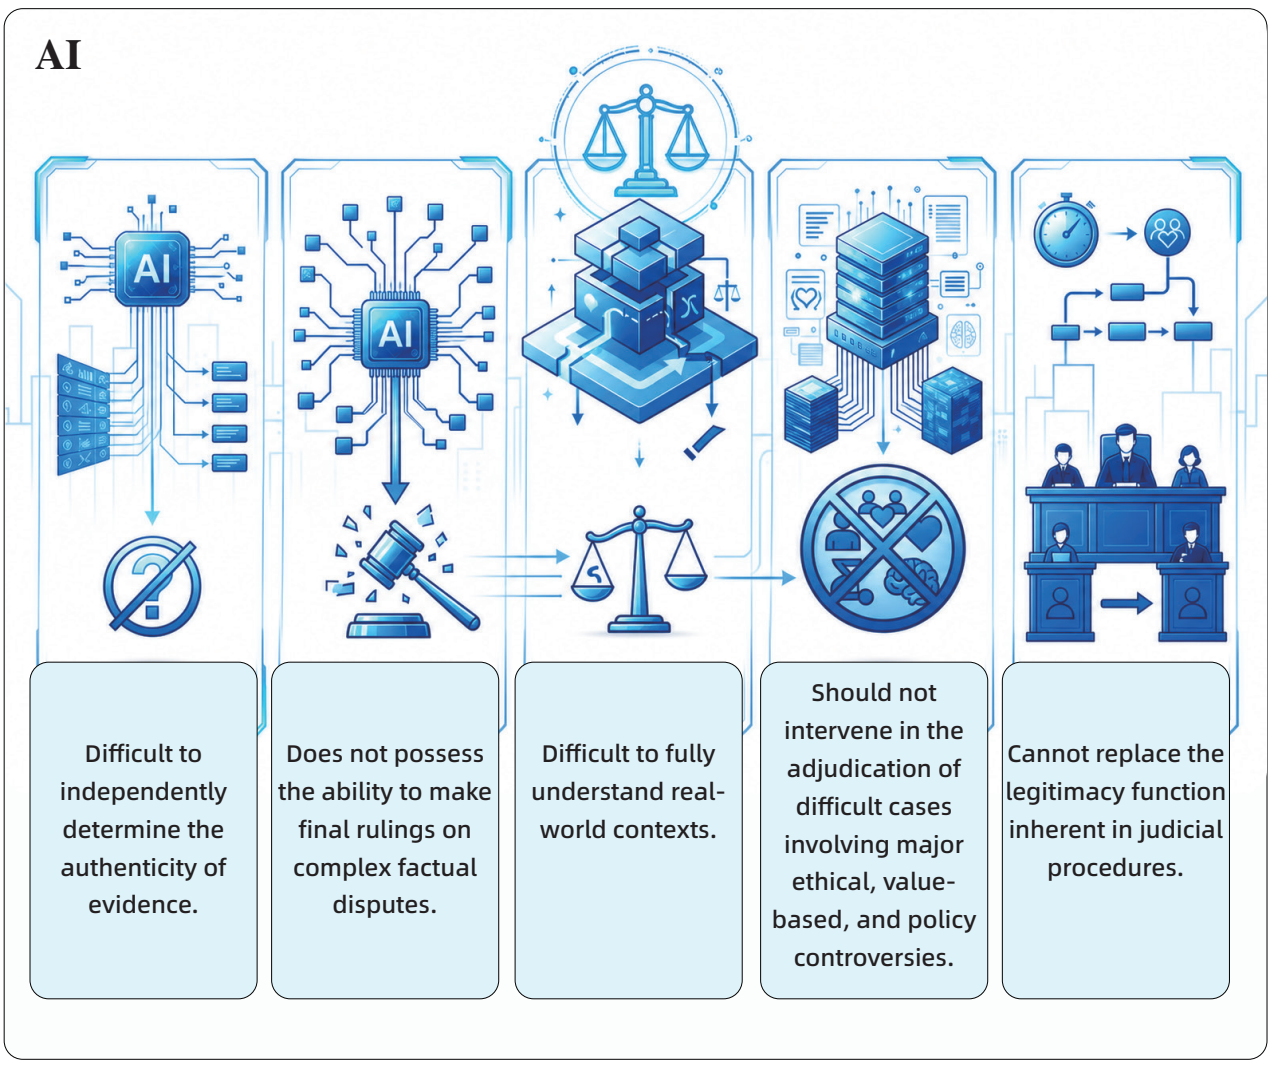
Creating a "digital-intelligence-native" mode of judicial work means that judicial personnel should possess the ability to collaborate with AI. Future judges, judicial support staff, clerks, and adjudication management personnel must not only be familiar with legal norms, litigation procedures, and adjudication work, but also understand the basic principles, application scenarios, boundaries of use, and risk characteristics of AI. Judicial personnel should develop the ability to use AI to rapidly read and understand various types of materials, including texts, images, audio, and video, and should be skilled in using intelligent tools to conduct similar-case searches, organize issues in dispute, check legal documents, prepare for hearings, and manage adjudication work. At the same time, judicial personnel should also possess the necessary capabilities for prompt design, output verification, risk identification, and bias correction, and be able to determine whether AI-generated content is accurate, complete, lawful, and appropriate. In areas where conditions are mature, key professionals may also become co-developers of AI tools, working with technical teams to sort out business rules, design application scenarios, label training data, and optimize model performance, thereby promoting the formation of intelligent products that truly meet judicial needs.

Creating a "digital-intelligence-native" mode of judicial work means that AI-assisted systems should be deeply embedded into adjudication processes, rather than remaining at the level of isolated, scattered, and superficial tool use. AI should play a role throughout the full lifecycle of cases, moving from assisting individual tasks to driving overall processes and promoting systematic transformation in case-handling methods. At the pre-litigation stage, it can assist with legal consultation, litigation risk alerts, dispute diversion, diversified dispute resolution, and preliminary review of materials. At the case-filing stage, it can assist with material recognition, cause-of-action identification, jurisdiction review, element extraction, and procedural diversion. At the trial stage, it can assist with case-file reading, summarization of issues in dispute, generation of hearing outlines, similar-case searches, prompts on application of law, drafting of legal documents, and document verification. At the enforcement stage, it can assist with analysis of property clues, identification of enforcement risks, recommendations on enforcement measures, and management of cases in which the current enforcement proceedings have been closed after unsuccessful enforcement. At the adjudication management stage, it can assist with case trend analysis, quality and efficiency evaluation, trial-limit warnings, monitoring of deviations in judgments, and organization of adjudicative rules for similar cases. Through full-process, systematic, and scenario-based application, AI can truly become an important component of court operational capacity. In line with this development trend, new capabilities such as full-process case-handling assistance agents, multi-agent collaboration, and continuous accumulation of judicial knowledge will gradually become important components of "digital-intelligence-native" courts.



Shenzhen Courts expect that, in the future, AI will be incorporated into the overall process of case handling, promoting a gradual shift in judicial work from a model of "human-led, AI-assisted" to a new model of "human-machine collaboration, intelligence-driven," thereby transforming how justice is produced. Judicial experience will be continuously accumulated as inherent AI capabilities, promoting the sustained transformation of judicial knowledge into productive forces. Multi-agent collaboration will become an important form of intelligent judicial work, proactively planning a large amount of judicial work and promoting the formation of a "digital-intelligence-native" mode of judicial work.

Of course, creating a "digital-intelligence-native" mode of judicial work does not mean interpreting AI capabilities as unlimited, still less does it mean handing judicial judgment over to AI. On the contrary, the deeper the application of AI, the more clearly we must understand the boundaries of its capabilities, and the more clearly we must define the bottom line of judicial assistance functions. At present, existing LLMs and related AI systems have already demonstrated strong capabilities in legal text comprehension, information retrieval, document generation, element extraction, and language expression. As technology continues to iterate, its capabilities will continue to improve. However, for a considerable period in the future, AI will still find it difficult to fully perform several core tasks in judicial work.



First, AI finds it difficult to independently determine the authenticity of evidence. Judicial proof activities involve not only the content of evidence itself, but also complex factors such as the source of evidence, the process of its formation, collection procedures, statements of the parties, witness credibility, cross-examination in court, and rules of experience. AI can help organize evidence, extract information, compare contradictions, and indicate connections, but it cannot replace judges in making a comprehensive judgment on the authenticity, legality, relevance, and probative force of evidence.

Second, AI does not possess the ability to make final rulings on complex factual disputes. Fact-finding often requires comprehensive weighing among incomplete, uncertain, or even conflicting evidence, and must take into account court investigation, the parties' production and cross-examination of evidence, the formation of the judge's inner conviction, and requirements concerning standards of proof. AI can assist judges in discovering factual clues, constructing factual timelines, and sorting out disputed issues, but it cannot replace judges in forming inner conviction, nor can it serve as the final subject of fact-finding.

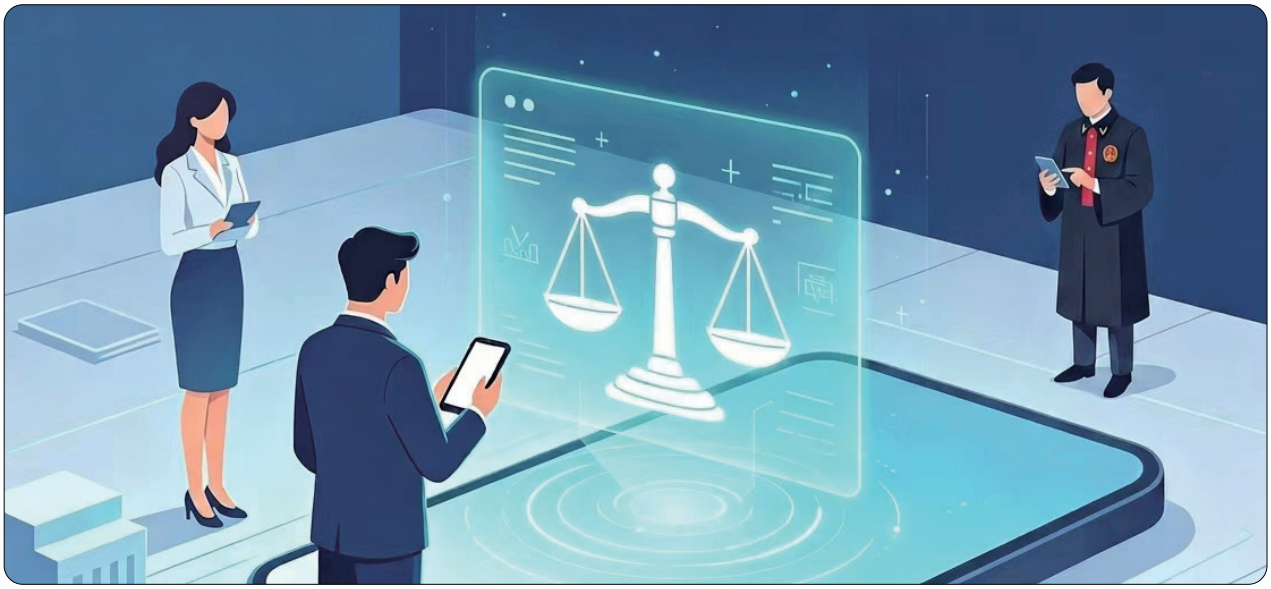
Third, AI still finds it difficult to fully understand real-world contexts. The application of many legal standards depends on comprehensive cognition of the physical world, social life, and industry practices, such as "reasonable duty of care," "obviously unfair," "serious misunderstanding," "obviously improper," "necessary limits," and "ordinary standards." These judgments are not merely matters of text matching; they also involve experience, context, common sense, social evaluation, and institutional purpose. At the present stage, AI still lacks a stable and reliable "world model," making it difficult for AI to independently complete legal evaluations that require deep real-world cognition.

Fourth, AI should not intervene in the adjudication of difficult cases involving major ethical, value-based, and policy controversies. Some cases may involve complex value balancing concerning human dignity, family ethics, public interests, social policy, protection of vulnerable groups, and conflicts between fundamental rights. The adjudication of such cases is not merely a matter of legal application; it also reflects the judicial organs' prudent judgment on fairness and justice, public order, and social values within the framework of the Constitution and laws. AI may provide support in materials retrieval and comparative research, but it should not generate or determine conclusions involving value judgments.

Fifth, AI cannot replace the legitimacy function inherent in judicial procedures. The authority of judicial decisions derives not only from the correctness of their conclusions, but also from the openness, fairness, participation, adversarial nature, and explainability of the procedures. Procedural stages such as court investigation, production and cross-examination of evidence, debate and statements, judicial clarification by judges, and collegial deliberation are important institutional arrangements through which parties participate in justice, express their views, and accept judgments. AI must not compress parties' procedural rights in the name of efficiency, nor may algorithmic judgment replace adjudication procedures conducted in accordance with the law.

Looking to the future, the judicial application of AI by courts should be based on the internal rules of judicial work and technological trends, with the goal of improving judicial quality, efficiency, and credibility. By empowering court work with AI, courts can enable the public to experience fairness and justice through judicial services that are more convenient, efficient, inclusive, and transparent.

6.3 Building a Smart and Convenient Online Litigation Service System



The development of online litigation is required by the transformation of judicial operations in the digital and intelligent era, and it also reflects the people's realistic expectations for judicial services that are more convenient, efficient, inclusive, and fair. As technologies such as the Internet, big data, and AI become deeply integrated into social life, profound changes have taken place in how people participate in litigation, how they access judicial services, and what they expect in terms of judicial efficiency and transparency. Online litigation should no longer be merely a simple migration of offline litigation activities onto online platforms. Rather, it should reshape how disputes are filed, managed, and resolved, optimizing the litigation experience, improving judicial efficiency and quality, and promoting substantive dispute resolution.

At present, Shenzhen Courts have largely completed the digitization of the full litigation process and have formed a relatively mature application foundation in areas such as online case filing, online payment, electronic service of documents, online mediation, online trial, electronic case files, electronic signatures and seals, and online document review. The next step is to further advance intelligentization on the basis of digitalization, empower online litigation services with AI, and promote the transformation of judicial services from being "available online" to being "easy to handle online," "intelligently handled," and "precisely handled." This is an important logical starting point for the future development of digital courts.

6.3.1 Promoting the Sharing of Judicial AI Services Across Society, and Enhancing the Convenience and Inclusiveness of Online Litigation Services

A digital court is, in essence, a "two-sided platform" connecting the public and the courts: one side connects the providers of judicial services, namely courts, judges, and other judicial personnel; the other side connects the users of judicial services, namely parties, litigation agents, and other litigation participants. An AI-assisted adjudication system can not only empower judges and internal court workflows, thereby improving the supply capacity of judicial services, but can also, through digital courts, push AI assistance to the earliest stage of the litigation process. In this way, people can directly experience the convenience and efficiency brought by AI technology throughout the entire litigation process, including case filing, evidence submission, cross-examination, trial, mediation, applications, service of documents, and case-file review.

Future digital courts should focus on the high-frequency needs, pain points, and operational barriers encountered by the public in online litigation, and develop online litigation guidance systems that are more intelligent, user-friendly, and easy to understand. Through functions such as intelligent Q&A, process navigation, generation of material checklists, risk alerts, and node reminders, the system can provide parties with litigation services that are human-centered, personalized, and precise. Online evidence submission, online cross-examination, online applications, and supplementary filings should be made more intelligent. For materials uploaded by parties, such as contracts, invoices, chat records, transfer vouchers, photographs, videos, and appraisal opinions, the system can assist with format recognition, content extraction, element labeling, chronological ordering, and evidence classification, guiding parties to submit more complete, clear, and effective evidence materials around the disputed issues. Where evidence materials are incomplete, unclear, repetitive, or have unclear relevance, courts may explore the use of AI to provide appropriate prompts, provided that such prompts do not affect the parties' status as litigation subjects or exceed the authority of the court. This can help parties improve the quality of their evidence submission and also help judges improve the efficiency of case-file review and adjudication.

Digital courts should explore providing parties with intelligent services that combine LLMs with searchable legal knowledge bases. By accessing reviewed, traceable, authoritative, and reliable laws and regulations, judicial interpretations, guiding cases, typical cases, adjudicative rules from similar cases, and litigation guidelines, such systems can provide the people with auxiliary services including legal basis searches, similar-case references, evidence organization, standardized expression of claims, and clarification of disputed issues. These services are public legal information tools, not substitutes for professional legal advice. They must not replace professional legal services provided by lawyers and others, must not make definite promises regarding case outcomes, and must not mislead parties into forming unreasonable expectations.

At the same time, the development of online litigation must attach great importance to the issue of the digital divide. The value of online litigation does not lie in the mere pursuit of technological advancement, but in enabling more people to access justice more fairly and conveniently. For elderly persons, persons with disabilities, persons with relatively low levels of education, and other litigant groups unfamiliar with digital technology, more adequate technical assistance and alternative services should be provided.

6.3.2 Enhancing Judicial Transparency and Judicial Credibility Through "Predictable Justice"

Judicial fairness requires not only that adjudicative outcomes comply with the law and the truth of the facts, but also that the process by which judgments are formed be open, transparent, participatory, and understandable. Litigation is not a closed process in which courts unilaterally exercise adjudicative power. Rather, it is an institutionalized process in which parties participate in accordance with the law, express their opinions, produce and cross-examine evidence, and engage in debate, and in which the court renders a judgment on the basis of procedural guarantees. The development of online litigation and AI technology provides new technological conditions for enhancing parties' participation in litigation, improving the transparency of the judicial process, and making adjudicative outcomes more understandable and predictable.

Predictable justice does not mean enabling parties to know the result of a case with certainty at the beginning of litigation. Rather, it means enabling parties, through sufficient information disclosure, procedural guidance, rule explanation, and interactive feedback, to understand how litigation procedures will proceed, which disputed issues the court will focus on, which facts need to be proved, which legal

norms may apply, and whether their opinions have been fully heard and carefully considered. By improving the understandability of the litigation process and the explainability of judgment formation, misunderstandings and confrontation caused by information asymmetry can be reduced, and parties' acceptance of and confidence in adjudicative outcomes can be strengthened.

During the online litigation process, parties can fill in, confirm, supplement, and submit case-related information and materials through the platform. On this basis, the digital court can intelligently identify and structurally process the unstructured materials submitted by parties, assisting in the generation of basic case facts, claims, defense opinions, disputed issues, evidence lists, legal bases, similar-case references, and other content. Digital courts can then feed the processed derivative data and preliminary summarized results back to the parties online for confirmation and solicit opinions from all sides. For cases involving complex facts, numerous disputed issues, or unclear statements by parties, AI agents may also be used to conduct multiple rounds of interaction, guiding parties to provide supplementary explanations around key facts, core disputes, and evidentiary issues.

The significance of this mechanism lies in the following:

First, it can simultaneously transform the process of parties' participation in litigation into a process of legal education and improvement of litigation capacity. Under the guidance of the system, parties can gradually understand claims, burdens of proof, rules of evidence, legal bases, and procedural requirements, which helps foster more rational and standardized litigation conduct.
Second, it can enhance parties' understanding of case facts and the application of law, enabling them to form more reasonable expectations regarding the direction of the case and the adjudicative outcome, and reducing dissatisfaction arising from unfamiliarity with procedures, unclear rules, or insufficient communication.

Third, through a "human-in-the-loop" confirmation mechanism, it can improve the accuracy of case data and information summaries. Fact summaries, disputed issues, and evidence organization results generated by the system must be subject to party feedback, judicial review, and procedural confirmation before they can be used as adjudication-assistance references. This helps eliminate information errors to the greatest extent possible and promotes fact-finding.

Fourth, it enables judges to gain a more comprehensive and accurate understanding of the full picture of a case before the hearing, improving the focus of hearings and the quality of judicial reasoning.
At the same time, AI assistance for judicial reasoning should be integrated with online legal explanation. AI can assist judges in checking whether the fact-finding in judgments is complete, whether responses to disputed issues are sufficient, whether citations of legal bases are accurate, whether the reasoning structure is clear, and whether adjudicative standards are consistent with those in similar cases. After a judgment is rendered, courts may explore the provision of services through online platforms, such as summaries of key points of the judgment, explanations of legal bases, reminders concerning performance obligations, and notifications of appeal rights, so as to help parties understand the content of the judgment and subsequent procedural arrangements. Through full-process, visualized, and interactive judicial services, courts can help ensure that the people not only "receive a judgment," but also "understand the judgment" and "trust the judgment."

6.3.3 Promoting the Mutual Reinforcement and Integration of Procedural Justice and Adjudication Efficiency Through Digital-Intelligence Empowerment

Empowering online litigation with digital-intelligence technologies is not only about increasing the speed of case processing. More importantly, it is about enabling judicial efficiency and procedural justice to reinforce and complement each other through process optimization, resource reallocation, and procedural strengthening. A truly modern judiciary responds faster, serves more precisely, safeguards procedure more fully, and delivers higher-quality judgments.

Digital courts should explore establishing a working mechanism of "distributed pre-trial preparation plus concentrated trial." "Distributed pre-trial preparation" means moving forward to the pre-trial stage a large amount of work that was previously concentrated at the trial itself, such as fact organization, presentation and cross-examination, summarization of disputed issues, procedural notifications, and mediation communication. With the help of online platforms and intelligent tools, such work can be advanced asynchronously, in a decentralized manner, and with traceable records. A "concentrated trial" means that, on the basis of sufficient pre-trial preparation, clearly defined disputed issues, fixed evidence materials, and clarified procedural matters, trials can focus on core disputes, key evidence, and issues concerning the application of law, thereby improving the substantive quality of the trial.

This mechanism helps promote a transformation in judges' case-handling methods from the traditional model of "finishing one case before moving on to the next" toward a model of "parallel processing of multiple cases, node-based advancement, and concentrated adjudication." With system support, judges can simultaneously advance pre-trial preparation for multiple cases, arrange trials and judgments according to the maturity of each case, substantially reduce waiting time and procedural idling, and improve the efficiency of judicial resource utilization. For parties, they can submit materials, express opinions, and participate in procedures online in a more flexible manner, reducing the time costs associated with traveling to and from court and waiting for scheduling. For courts as a whole, this can optimize the allocation of adjudication resources, auxiliary tasks, and technological capabilities, thereby improving case turnover efficiency and judicial supply capacity.

More importantly, the more thorough the pre-hearing preparation, the clearer the focus of the hearing, and the fuller the procedural participation, the more parties can feel that their opinions have been carefully heard, their claims incorporated into the adjudication framework, their evidence fully examined, and the judgment made on the basis of full expression and dialogue among all parties. Such a litigation process not only helps judges ascertain facts and apply the law more accurately, but also helps improve parties' acceptance of judgment, and helps close cases in a way that truly resolves the underlying dispute.

6.4 Expanding the Public Service and Social Governance Functions of Judicial AI

Courts are an important component of the national governance system. Through adjudication, the judiciary settles disputes, confirms rules, evaluates conduct, and restores order. It not only resolves individual disputes, but also communicates legal norms and value orientations to society. The value of judicial AI services should not be limited to internal adjudication and enforcement activities within courts, nor should it serve only the handling of individual cases. Over time, the judicial process has accumulated a vast amount of legal knowledge, adjudicative rules, similar-case experience, dispute patterns, and governance information. Such knowledge and data can not only help judges and judicial support personnel handle cases more effectively and help parties participate in litigation more effectively, but can also, on the premise of legality, compliance, security, and controllability, be extended to fields such as public legal services, law-related publicity and education, administrative law enforcement supervision, social risk warning, and grassroots governance. In this way, judicial AI can become an important resource for improving the standard and reach of rule-of-law governance.

6.4.1 Judicial AI Services Can Improve the Public Legal Supply and Promote Inclusive Sharing of Legal Knowledge

The legal knowledge, adjudicative rules, and case-handling experience accumulated in the judicial process are highly professional, practical, and authoritative. AI models, knowledge bases, and intelligent agents formed through continuous accumulation, verification, and refinement in real judicial work may, within an appropriate scope, be used in public legal service scenarios such as legal publicity, auxiliary legal consultation, litigation risk alerts, and dispute resolution guidance, helping the public access basic legal knowledge and behavioral guidance more conveniently.

Efforts should be made to explore the development of a public-facing judicial knowledge service platform, effectively lowering the threshold for the people to obtain legal information and improving the coverage and precision of public legal services. Based on authoritative resources such as laws and regulations, judicial interpretations, typical cases, adjudicative rules from similar cases, litigation guides, and Q&A on common disputes, judicial AI can provide services including intelligent Q&A, case retrieval, risk alerts, explanations of litigation procedures, guidance on evidence preparation, and mediation suggestions. For high-frequency disputes relating to people's livelihoods, the system can explain relevant legal rules, common disputed issues, allocation of the burden of proof, available dispute resolution paths, and points requiring attention in litigation in plain and accessible language, helping the public form rational judgments at an early stage of a dispute.

6.4.2 Judicial AI Services Can Promote Source Prevention and Diversified Resolution of Conflicts and Disputes, Facilitating Dispute Resolution Before Litigation

Judicial AI applications can serve not only cases that have already entered the courts, but also potential disputes that have not yet entered litigation procedures. Through intelligent analysis of common types of disputes, adjudicative rules, mediation experience, and litigation risks, it can provide auxiliary support for people's mediation, administrative mediation, industry mediation, commercial mediation, arbitration, notarization, grassroots governance organizations, and other mechanisms, promoting the resolution of more conflicts and disputes at the grassroots level, before litigation, and at the earliest possible stage. AI can also provide mediators with support such as similar-case references, rule prompts, auxiliary generation of mediation proposals, and performance risk alerts. For cases suitable for mediation, where the facts are relatively clear, disputed issues relatively defined, and legal consequences relatively predictable, judicial AI can guide parties to rationally assess litigation costs, time costs, and performance risks, thereby promoting the formation of mediation proposals that are more targeted and enforceable.

6.4.3 A "Digital-Intelligence-Native" Judicial Process Can Support Social Governance and Policymaking Through Enhanced Data Analysis and Feedback

Courts are not isolated adjudicative bodies; their work is intertwined with the wider economic and social environment and in the national governance system. Judicial data provides a concentrated reflection of conflicts and disputes, risk factors, and governance shortcomings in economic and social operations. By collecting, cleaning, analyzing, and assessing judicial data in accordance with the law and compliance requirements, courts can identify emerging and trend-based issues in specific fields, industries, and regions, thereby providing references for relevant departments in carrying out social governance.

Relying on judicial data, governments can conduct analysis and assessment regarding areas with high incidence of similar cases, risks in key industries, and difficulties in grassroots governance, so as to promptly identify the causes of frequent disputes, especially deficiencies in institutional supply. At the same time, judicial AI services can also play an auxiliary role in standardizing administrative law enforcement, improving public policies, and supporting scientific legislation.



Conclusion

Bringing Technology Closer to People and Making Justice More Accessible

The entry of AI into the judiciary is not intended to make justice cold, nor is it intended to hand adjudication over to machines. On the contrary, its true value lies in freeing judges from a large amount of repetitive and routine work, allowing them to spend more time listening to statements, ascertaining facts more deeply, explaining the law more fully, and making judgments more prudently. It also lies in enabling the public to make one fewer trip, wait less, face less uncertainty, and gain more convenience, more understanding, and greater trust in fairness and justice.

Shenzhen Courts look forward to and will continue to promote the construction of a "digital-intelligence-native" court form. We believe this is the future of judicial work. In Shenzhen's practice of applying AI in the judiciary, it is foreseeable that a "digital-intelligence-native" court will not merely optimize litigation processes and judges' case-handling workflows, but may also profoundly transform the organizational methods of judicial procedures, the allocation of judicial resources, and the mechanisms through which judicial judgments are formed. Facing this future, we need to begin thinking together and continue exploring from today onward. We also sincerely look forward to the participation of more people of insight, so that we may jointly promote the innovative development of judicial work in the era of AI.

Appendix

Appendix 1: Guide for Using AI Agents in the Judicial Domain V2.0 (Sept 2026)
Appendix 2: Glossary
Appendix 3: Chronology of Major Events

Appendix 1

Guide for Using AI Agents in the Judicial Domain V2.0 (Sept 2026)

One-Page Overview – Building a governance framework of prudent integration, strict risk prevention, and practical mastery.

Why Regulation Is Needed: Four Kinds of New Risks	Authority Overreach:	Task Contamination:	Process Deviation:	Loss of Control:
	Core Judicial PowersBeing "Hollowed Out"	Task Objectives MaliciouslyDistorted Through "Autonomy"	Dynamic Execution Chains Going Off Track	Privilege Abuse and Cascading Failures

Four Positive Principles

Prudent Integration: Upholding the Core Principles for Using AI Agents in the Judiciary

Positioning AI only as Assisting Tools
Keep the judge's decision in the loop, uphold the Principle of Immediacy,and never allow AI to replace judges in adjudication in any way.

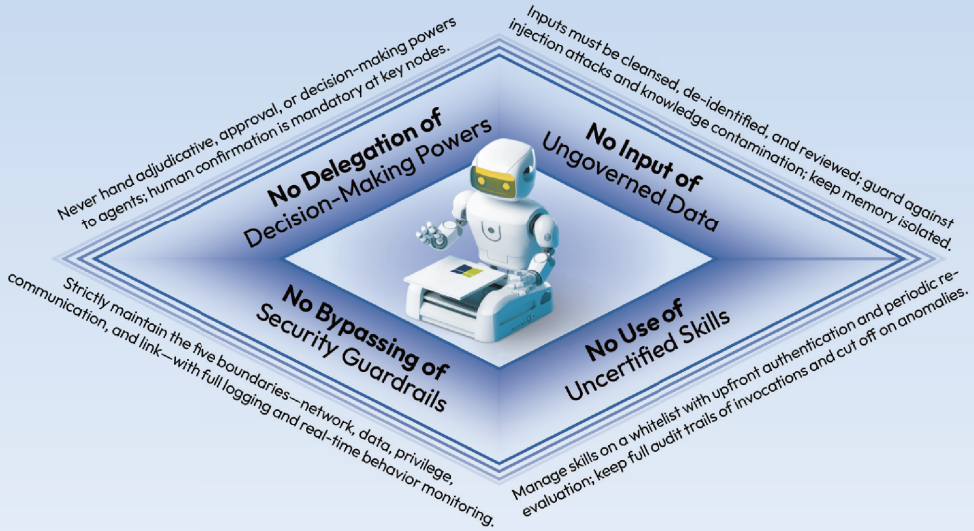
Deeply Integrate AI with Adjudicating Process
Fit the rules of adjudication and judges' working habits; deeply integrate the multi-model engine, business systems, and authoritative knowledge bases.

Implementing judicial accountability synchronously
Let those who hear the case adjudicate and those who adjudicate be accountable; pre-set review and confirmation nodes; record the decision-making process as it happens.

Data security throughout the entire process (3.4)
Deploy on a dedicated network in a closed environment; strictly observe the data-security red line across the full life cycle of training, inference, and application.

Four-Point Negative List

Guarding Against Risk: Defining the Negative List for Judicial Use of AI Agents



Four-Stage Practical Guidance

Practical Mastery: A Step-by-Step Guide for Judicial Users

Define the Task
Goal + context + skills + output + constraints — all five elements in place.
Define it precisely

Supervise the Process
Check execution steps, skill calls, cited sources, and any skipped steps.
See through it clearly

Verify the Output
Verify facts, statutory grounds, cases, logic,and caliber, with tiered, evidence-based checks.
Verify it rigorously

Harness the Loop
Correct, tighten, calibrate, and supplement; iterateto convergence and accumulate experience.
Refine it diligently

Run — Verify — Correct — Rerun, converging step by step.

1. Introduction

This guide provides practical guidance for judges and judicial staff on using AI agents and serves as a reference for judicial AI development. It emphasizes that the use of AI agents should remain auxiliary, standardized, secure, and consistent with accountability requirements. The goal is to integrate AI deeply into judicial work–improving the quality and efficiency of adjudication while upholding judicial ethics, data security, and justice as a whole.

An AI agent is an intelligent system that can perceive its environment, make decisions, and take actions to achieve a specific goal. In the judicial context, this means AI applications designed to support specific judicial tasks. With recent breakthroughs in autonomous agents, these systems can now plan tasks, reason step by step, invoke skills, and interact with their environment–powered by Large Language Models (LLMs), agentic workflows, and interleaved thinking. Their "autonomy" lies in the ability to dynamically formulate strategies and steps to accomplish a task.

2. The Future Is Here: Opportunities and Challenges of AI Agents in the Judiciary

Since 2026, agent-based AI applications have made significant progress, bringing new opportunities to the judiciary in three ways. First, they boost efficiency by automating repetitive tasks such as document service, data entry, and routine inquiries. Second, they respond flexibly to individualized case management and litigant service needs at lower cost. Third, they lower the development barrier–with modular, configurable, and pre-defined agents, complex judicial AI applications become easier to build and deploy.

At the same time, AI agents introduce serious risks due to their high degree of autonomy and interactivity. Without proper controls, these risks can undermine judicial procedures, accountability, and public trust. The main risks are:

2.1 Authority Overreach: Core Judicial Powers Being "Hollowed Out"

When AI agents are used, judicial powers may be improperly transferred. Judges may over-delegate—either because they overestimate what agents can do or because they become reliant on the convenience. For example:

Delegation: Giving only the final instruction (e.g., generate a judgment) while skipping the intermediate steps of fact-finding and legal reasoning, turning the decision-making process into a black box
Wholesale outsourcing: Handing over core judicial work, such as evidence review and fact determination, to the AI agent without differentiation, even though these require the personal involvement of judges. This in effect creates an automated adjudication chain that undermines judicial accountability.

2.2 Task Contamination: Task Objectives Maliciously Distorted Through "Autonomy"

Instructions given to AI agents can be misinterpreted or tampered with during complex execution. An autonomy depends on "informational autonomy" at intermediate steps. This means the set of instructions it receives, the context it reads, and the return values from skill invocations can all be contaminated or altered in transit, allowing the real task objective to be maliciously rewritten. For example:

User-initiated attacks: Users, including litigants acting through proxy tools or textual materials, may try to bypass the agent's built-in safety constraints through role-playing, encoding bypasses, multi-turn gradual inducement, or system prompt leakage. In judicial settings, particular vigilance is needed against attacks conducted through submitted case materials, leading questions, or personalized parameter settings.
Third-party prompt injection: Malicious instructions may be embedded in third-party content, tricking the agent into reading and executing them. These may be steganographically hidden in images, audio, or scanned documents to bypass text-based filters.

Knowledge and memory contamination: Litigation materials may contain AI-generated fabricated legal provisions or fictitious cases. If this "contaminated knowledge" enters the task context unchecked, all subsequent reasoning will rest on false premises. For agents with long-term memory, erroneous or malicious content injected in one task can persist across sessions and affect all future tasks, creating what amounts to "chronic contamination."

2.3 Process Deviation: Dynamic Execution Chains Going Off Track

Skills and self-generated execution steps that agents dynamically invoke may deviate from procedures or internal regulations. For example:

Unverifiable steps: When AI agents autonomously invoke external skills or create their own tools, the process of generating intermediate steps may lack transparency and explainability, making compliance difficult to audit.

Unauthorized skill introduction: A skill invoked or created by the AI agent may contain faulty logic, malicious code, or provisions that violate judicial procedures. The task may appear complete while actually breaching case-handling norms. This risk is especially acute in the capability component supply chain: bundled extensions, third-party tool interfaces, code-generated components, and dynamically spawned subagents can all serve as contamination entry points. Malicious capability components may bypass runtime security policies entirely, operating outside the kernel, where traditional security reviews cannot cover them.

2.4 Loss of Control: Privilege Abuse and Cascading Failures

Once AI agents gain system-level privileges, they may trigger unauthorized operations, data leaks, and other security incidents. For example:

High-privilege misoperation: With unclear privilege boundaries and insufficient anomaly detection, agents may invoke high-privilege interfaces or trigger critical business processes by mistake, causing irreversible interference with case adjudication.Data boundary breaches: AI agents may attempt to access sensitive data beyond their authorized scope to complete a task, creating data leakage risks.Multi-agent cascading errors: In multi-agent collaborative scenarios, a single AI agent's anomalous operation can trigger chain-reaction errors through task dispatch and cross-component invocation chains, amplifying the scope of impact.

3. Prudent Integration: Upholding the Core Principles for Using AI Agents in the Judiciary

3.1 Positioning AI only as Assisting Tools

Regardless of how the technology evolves, the judge's autonomous decision-making authority should be respected. The Principle of Immediacy (亲历性原则 , a foundational doctrine in Chinese judicial procedure requiring judges to directly participate in adjudication) must be upheld throughout and let the people who directly hear the case make rulings and other core decisions. In the development and use of AI applications, the judge's decision must always remain in the loop of AI use.

3.2 Deeply Integrate AI with Adjudicating Process

Incorporating AI must be aligned with how adjudication actually works and how judges handle cases. Harness engineering combines multi-model engine integrating general-purpose and domain-specific judicial engineering, and authoritative legal knowledge bases. Tailored to judicial AI components, it pinpoints the hardest tasks at critical case-handling stages, designs around each case's standards, and delivers results that meet actual needs.

3.3 Implementing judicial accountability synchronously

Regardless of how the technology evolves, accountability stays with the judge. The AI applications must include mandatory checkpoints for judge review, confirmation, and decision-making to ensure judgments are made by the adjudicator and judicial accountability rests with the adjudicator. AI may provide decision support, knowledge push, and drafting assistance, but the decision-making process must be recorded to ensure the principle that those who hear cases shall make the judgments, and those who make the judgments shall be held accountable.

3.4 Data security throughout the entire process.

Judicial data contains substantial private information and represents an inviolable security red line. AI data must be protected across its entire lifecycle. AI must be deployed on dedicated networks, managed by dedicated personnel, with dedicated security architecture. All judicial data, whether used for training, fine-tuning, inference, or application, must remain within an isolated environment, eliminating the risk of sensitive data leakage at the root.

4. Guarding Against Risk: Defining the Negative List for Judicial Use of AI Agents

With the latest advances in technology, LLM-powered agents raise the capability ceiling but also introduce new security risks.Meanwhile, excessive constraints may also limit the performance of agent technologies and hinder the unleashing of productivity. To strike a balance between application potential and security risks, and on the basis of complying with the four aforementioned principles, the following" four-negative list" is proposed. It defines the security baseline and application red lines through explicit prohibited behaviors, establishing clear areas of exclusion while preserving reasonable room for beneficial innovation.

4.1 No Delegation of Decision-Making Powers

Under no circumstances may the power to adjudicate, approve, or determine a judicial task be delegated to AI agents. The fundamental model is "judges decide, agents execute." All substantive evidence admission, fact finding, rule application, and adjudicative conclusions must be made explicitly by judges. Technical measures must prevent agents from offering substitute decisions in the form of "adjudication recommendations" or "one-click generation." The system should intercept such instructions. In multi-step agentic workflow orchestration, all critical decision checkpoints requiring "judges in the loop" confirmation must be retained and enforced in real-world use. Automated chains that bypass human judgment shall be strictly prohibited to avoid any form of black-box or misleading decision-making.

4.2 No Input of Ungoverned Data

Raw data that has not been cleansed, desensitized, and security-reviewed must not be used to train agents or injected into their task context. All inputs, including prompts, case materials, and knowledge, must go through full-process governance: source verification, content cleansing, knowledge validation, and sensitive information handling. Effective defenses against prompt injection attacks are mandatory. In principle, raw litigation materials that have not been structurally extracted must not be used directly as agent inputs, cutting off the injection pathway at the source. Legal knowledge for reasoning must come from certified authoritative databases, and potentially fabricated legal information in litigation materials must be flagged. Context retrieved by agents through external skills, including web search results, must undergo instruction filtering and format isolation. User inputs must be screened for jailbreak attempts.

Long-term memory, cross-session state, and decision caches must be periodically cleansed, isolated by case, and cleared by task to prevent a single attack from spreading across sessions into chronic contamination.

4.3 No Use of Uncertified Skills

Agents must not invoke external skills, tools, plugins, or other capability components that have not undergone security assessment, functional certification, and compliance review. All available skills must be governed by a strict allowlist. Users must not install or configure uncertified third-party skills which are not on the allowlist. Technical measures must also be taken to prevent AI agents from creating or acquiring new skills on their own, and constrain them to the pre-approved toolset. All AI and capability components must pass integrity checks before loading; invoked capability components and activated subagents must be verified. A unified skill management mechanism, jointly overseen by technology and knowledge-management departments within the judiciary, must handle pre-deployment certification, compliance review, and periodic post-deployment re-evaluation, with change approval and rollback procedures. All skill invocations must be fully logged into an auditable trace chain. Real-time monitoring and anomaly detection must be in place: if a capability component invocation deviates from expected behavior or triggers a security guardrail, the system should automatically interrupt the task chain, roll back to a safe state, and notify the adjudicating authority.

4.4 No Bypassing of Security Guardrails

All guardrails, including network isolation, access controls, data boundaries, audit logs, and circuit breakers built into the system, must be strictly maintained. Among them, the most critical are five categories that must not be bypassed or undermined.Network boundary. AI must run in strictly isolated network environments. Unauthorized access to external networks or cross-domain communication is strictly prohibited.Data boundary. Strict data classification and least-privilege access controls should be enforced. Agents may only access desensitized data that is necessary for the current task and within the operator's scope of authorization.Privilege boundary. The privileges of an AI must be tied to user identity, following "case-specific authorization" and "role-based authorization" principles. Each task runs in an independent sandbox to prevent privilege abuse and lateral movement.Communication boundary. In multi-agent scenarios, inter-agent messages must travel through certified encrypted channels. Unauthenticated communication or direct interaction bypassing middleware is prohibited. Protocol version controls must prevent downgrades to insecure protocols.Chain boundary. To address cascading failure risks in multi-agent task chains, task-level and system-level circuit breakers must be in place. When a single agent exhibits anomalous behavior (format errors, low confidence, timeouts), the entire downstream task chain is automatically suspended, the fault node isolated, and personnel notified. Throughout the whole process, all agent operations shall be fully logged, with real-time behavioral monitoring, anomaly alerts, and automatic-blocking mechanisms in place.

5. Practical Mastery: A Step-by-Step Guide for Judicial Users



As AI agents integrate more deeply into judicial workflows, they also raise the bar for users. The judiciary is a "high-stakes, tight-controls, precise-outputs" domain. Beyond the baseline requirement of using only securely deployed, officially certified tools, whether users can effectively harness the use of AI agents has become a critical competency for justice and efficiency. A clear trend is emerging: effective users are empowered, misusers are hindered, and skilled users achieve mastery. Establishing normative guidelines for judicial users is both a practical necessity and a fundamental requirement for the proper exercise of judicial power.

The following step-by-step guide, written from the user's perspective, organizes the use of AI agents into four stages: Define the Task, Supervise the Process, Verify the Output, and Harness the Loop. It helps judicial users harness agent capabilities and translate AI productivity into just, efficient, and safe outcomes, which will ultimately contribute to access to justice.

5.1 Define the Task

An AI agent can only execute effectively if the task is clearly defined. Precision starts with the task definition, which includes the prompt and relevant content. A vague instruction may cause the agent to skip legal knowledge validation, bypass the judge's decision checkpoints, or reference unreviewed information on its own, leading to accumulated deviations, inaccurate results, and misplaced reliance. Users should address five areas: task type selection, content specification, input preparation, constraint setting, and output format.

5.1.1 Choose the task type.

Matters involving core judicial accountability are generally pre-configured by system developers as predefined tasks or agentic workflows, ensuring the pipeline is standardized, transparent, explainable,

and accountable. Users should default to these and follow their built-in steps for material screening, information verification, knowledge validation, decision injection, and output verification. When using custom prompts to inject or modify core accountability matters, users must exercise extra care across all four stages (Define, Supervise, Verify, and Loop), and results may only be incorporated into core accountability processes after rigorous manual review.

5.1.2 Specify the content.

Users conduct legal research, information synthesis, and in-depth investigations by defining custom prompts. The specific details of the prompt will directly affect the quality of the outcome. A well-defined task includes: a clear objective, a defined context scope, which skills to use, the expected output requirement, and necessary constraints. The more specific the task content, the better the agent's planning. Vague prompts leave too much room for autonomous planning, increasing the risk of deviation and errors.

5.1.3 Prepare the inputs.

The completeness and accuracy of the context directly affect whether the agent's reasoning is sound and focused. Users should manage context deliberately: provide complete, accurate information while removing irrelevant noise. Highlight key information through targeted data entry, case material selection, and knowledge retrieval instructions. Avoid information overload, missing critical premises, and context contamination, all of which may lead to accumulated deviations and unwarranted assumptions.

5.1.4 Set reasonable constraints.

System developers have already built in certain safeguards against hallucination and over-authorization. But users should also know how to set boundaries. Explicitly prohibiting certain agent behaviors within the task is just as important for the credibility of the results. Clear constraints bound the agent's action space, reduce unauthorized behavior, limit deviations, and lower verification costs.

5.1.5 Prescribe the task output.

Task outputs generally need to follow a fixed format for downstream verification and use. Users can require the agent to produce output in a specific structure–for example, listing legal provisions cited, annotating information sources, presenting reasoning step by step, or creating tables. This makes verification easier, helps spot deviations, and ensures the output integrates smoothly into documents and subsequent workflows.

Example 1 (Good Practice):The task is defined as "Based on the attached court hearing transcript and evidence list (context), compile a list of disputed issues in the case (objective). For each disputed issue, itemize the points where the parties disagreed during the hearing, cross-reference each against the uploaded evidence and invoke skills to retrieve relevant legal provisions and map them correspondingly to each dispute point (skills). Note: the legal facts of this case occurred before the Civil Code (民法典 , which took effect on January 1, 2021, replacing prior civil legislation) came into force. Retrieve the laws in effect at that time, specify the version, and trace back to authoritative databases for easy verification (constraints). Present the results in tabular format (expected output)."

Example 2 (Lesson Learned): The task is simply defined as "Write the final judgment for this case." This task has a vague objective, no context, no decisions specified, no constraints, and no format requirements. In the Shenzhen AI-assisted adjudication system, this prompt will be rejected as delegating decision-making authority.Even if it is indeed an instruction from a judge, the system cannot allow the agent to replace the

judge's decision-making process with a simple task. This kind of usage also forces the model to fabricate context or introduce hallucinations to complete the task, increasing risks of misinformation, flawed reasoning, and black-box processing.

5.2 Supervise the Process

A well-defined task still needs precise execution and supervision. Unlike traditional "input-output" tools, agents understand the task, break it into steps, invoke skills and tools, generate intermediate results, and complete reasoning progressively, making the process visible and traceable. Each step's output can influence the next step's direction. If users only look at the final output and ignore the intermediate process, they cannot catch errors in time, and deviations get unconsciously adopted. Users need to understand how autonomous agents work, why information alignment and checkpoint intervention matter, and how to monitor task execution.

5.2.1 Ensure task alignment.

Agents typically display visible execution records: task steps, material reads, skill invocations, and source citations. Since accurate transmission of task content to the agent's context window is a prerequisite for correct execution, users should first verify full alignment. Check (among others): Have all instructions been injected without omissions? Has the context (knowledge texts, case materials) been properly loaded? Has the agent correctly understood the task and followed it in its subsequent planning?

5.2.2 Monitor the execution.

Agents execute by planning autonomously, invoking skills and tools as needed, and progressing step by step. Users should review intermediate steps closely. Watch for (among others): skipped steps, unwarranted assumptions, missing elements, and logical errors in the step-by-step plan, reasoning, skill invocations, and intermediate outputs. Also check whether the skills and data sources being used fall within the authorized scope.

5.2.3 Intervene at anomalous nodes.

Agent tasks generally allow users to pause and intervene. Users need to learn when and how to step in. Intervene immediately when conditions may affect the task's substantive results. Key situations include: critical context not loaded (risking hallucinations); incorrect legal premises introduced (risking logical confusion); skill invocation failures (risking process deviation); attempts to access unauthorized resources; and intermediate reasoning that diverges from the judge's decision.

Example 3 (Good Practice): After asking the agent to organize factual disputes and link evidence for a sales contract dispute case, the judge did not jump to conclusions. Instead, the judge first checked the execution records: (a) material reads–the pleadings, hearing transcript, and both parties' evidence lists were all loaded without omissions; (b) step plan–the agent decomposed the task into four steps consistent with the instructions; (c) skill invocations–evidence analysis used the built-in case file reader, not external sources of unknown origin; (d) intermediate results–the factual dispute summaries did not add facts beyond the materials. Only after confirming alignment at each stage did the judge proceed to output verification.

Example 4 (Lesson Learned): The agent had completed its analysis with critical materials missing. The judge paused the task, uploaded the missing evidence, and reran. The output then fully covered both parties' claims. The root cause was missing input, a problem detectable during execution, not at the output stage.

5.3 Verify the Output

The outputs of AI may become reference material for reasoning or foundational content for legal reasoning in judgments. Evidence-based verification is therefore mandatory. Although agent outputs are typically fluent, well-structured, and stated with confidence, looking professional does not equal accuracy. Users must apply the strictest standards to guard against hallucination, factual errors, and prompt deviation. Verification should be tiered and evidence-based, covering three dimensions: legal basis, factual statements, and output logic.

5.3.1 Verify legal basis.

Even when an AI has already connected to authoritative databases and attached source links, manual verification is still required. Cross-check every cited provision, judicial interpretation, and guiding case (指导案例 , selected cases published by China's Supreme People's Court as authoritative reference for lower courts, distinct from binding precedent in common law systems) against authoritative databases (such as the National Laws and Regulations Database). Pay special attention to whether cited provisions are still in force. Legal bases that the agent "proactively" cited, meaning those not explicitly requested in the prompt, require extra scrutiny, and the judge must personally determine whether they may serve as valid grounds.

5.3.2 Verify factual statements.

Although the reference information in AI agent outputs has typically undergone security governance and corroboration, governance cannot cover every risk. Compare fact summaries, disputed-issue compilations, and similar content against the original case file materials. Confirm the agent has not omitted key facts, confused the parties' claims, or added information not present in the source materials.

5.3.3 Verify output logic.

Even when legal bases and facts are both correct, the agent's analytical process may still contain logical leaps, misinterpretations, or hidden assumptions. Check (among others): Are the cited facts and legal bases accurate and traceable to the source? Are the analytical conclusions consistent with the judge's decisions? Do the conclusions genuinely follow from the premises? Is there concept substitution or overgeneralization? Is the analysis of different viewpoints balanced?

5.3.4 Not all outputs need the same level of scrutiny.

For procedural tasks like date calculations, formatting, and typo checks, a quick confirmation suffices. For substantive content involving legal application, fact-finding, and judgment reasoning, each item must be verified against original materials and authoritative sources through evidence-based review.

Example 5 (Good Practice): After the agent completed a legal research report on a contract termination dispute, the judge verified it in three layers. First, legal basis: opened every authoritative database link attached to cited provisions and judicial interpretations, checking content and current validity. Second, factual statements: compared the report's case facts against the original file, confirming nothing was omitted, confused, or added. Third, output logic: checked whether conclusions followed from premises and were consistent with the judge's decisions. Procedural items like deadline calculations got a quick check; sections intended for legal reasoning were traced item by item. Only after verification was the report used as reference material.

Example 6 (Lesson Learned): The agent completed an analysis of the defendant's defense submissions. The output was well-structured and fluent, citing "Article 563 of the Civil Code" (《民法典》, the principal civil law statute in China) with accompanying text. When the judge traced each citation, they found the paraphrased "provision text" did not match the original in the authoritative database. It had incorporated language from the now-superseded Contract Law (《合同法》,replaced by the Civil Code in 2021). Source tracing showed it came from the defendant's brief. The judge issued a targeted correction: "The quoted text in Section 2 does not match the original of Article 563 of the Civil Code. Correct it against the authoritative database and attach the source link."

5.4 Harness the Loop

In judicial scenarios, agent outputs rarely satisfy every precision requirement on the first pass. Multi-round iteration is not a sign of failure; it is the normal way for agents. Users converge initial results toward precise outputs through a loop of run, verify, correct, rerun.The loop works on two levels: within a single task, through targeted corrections, task redefinition, and output calibration; across tasks, through accumulated experience. Users should master five loop methods:

5.4.1 Point out the exact location, what is wrong, and what the fix should be. For example: "The reference to Article 580 in paragraph 3 is incorrect; it should be Article 577. Please correct." Focus each correction round on one main issue. Adjusting multiple things at once makes it hard to tell what worked. The more specific the correction, the more reliable the fix. Correct issues as you find them; do not wait for the full task to complete.

5.4.2 Redefine the task.

If the output direction is right but lacks detail, add follow-up context or instructions. If it has fundamentally gone off track or critical premises are missing, do not keep patching. Redefine the task from scratch, incorporating discovered issues as constraints. For example: "Do not search broadly; focus only on the contract-related provisions of the Civil Code." If the materials are too noisy, supplement or reselect them and rerun. For example, upload new evidence or a recent higher-court ruling before regenerating.

5.4.3 Calibrate output standards.

If the result is generally usable but the scope, format, or level of detail does not match requirements, calibrate by specifying explicit standards. For example: "Reorganize the issues according to this court's retrial review standards" (再审审查, a specific judicial procedure in Chinese courts for reviewing previously decided cases). Further refine the tone, format, and granularity as needed.

5.4.4 Know when to stop.

The iteration loop needs clear stopping conditions to avoid unlimited verification effort. Stop and transfer to manual processing when: (a) the same error recurs after two rounds of targeted correction, meaning the task has exceeded the agent's current capability; (b) the task involves sensitive, irreversible, or high-impact matters, manual review is mandatory regardless of results. Before stopping, you may ask the agent for a self-check (e.g., "verify the current validity of every legal provision cited in this document"), but the agent's self-check cannot substitute for the user's final verification.

5.4.5 Save effective task definitions, constraints, and verification checkpoints as personal or departmental prompt templates for reuse in similar tasks.

This continuously lowers the cost of defining tasks and improves first-round output quality. Experience templates should capture methodology, not individual case information—case details must never leak into other cases. Where agent memory is involved, strictly follow the requirements of case-level isolation and task-level clearing.

Example 7 (Harnessing the loop): A judge asked the AI agent to search for adjudication standards for a type of retrial case (再审 , a judicial procedure in Chinese courts for reviewing and potentially correcting previously finalized judgments). The first-round output included unrelated case categories; the judge narrowed the scope and reran. The second-round output still cited outdated standards that had been superseded by a new judicial interpretation; the judge pointed out the error and required reorganization based solely on currently effective regulations. The third-round output repeated the same type of error. The same error recurring after two rounds of targeted correction means the task exceeds the agent's current capability. Further iteration would only waste verification effort. The judge stopped the loop, completed the search manually, and reported the issue to the system administration department for continuous improvement.



Appendix 2

Glossary

Large Language Model (LLM):
A large-scale model trained on massive text data, with natural language understanding and generation as its core capabilities. It can perform a wide range of language tasks, including text generation, summarization, translation, question answering, and reasoning. Large language models are the core underlying technology for current AI agents.

Hallucination:
The phenomenon where a large language model generates content that appears plausible but is unreliable–such as fabricating nonexistent legal provisions, cases, literature, or facts. Hallucination arises from the probabilistic nature of how large language models generate text and cannot be entirely eliminated. In judicial settings, hallucination is an inherent risk that must be a key focus during output verification.

Agent:
An intelligent system capable of autonomously perceiving its environment, making decisions, and taking actions to achieve specific goals. Agents typically possess basic capabilities such as memory, planning, and tool use, and represent an important form of AI products and services. In the judicial context, the term refers specifically to intelligent modules or systems designed to perform specific adjudication-support tasks.

Agentic Workflow:
A working mode in which the task execution process is broken down into multiple predefined steps, with a large language model completing each step sequentially according to a fixed pipeline. In an agentic workflow, the task decomposition logic and step sequence are preset by the system developer. The model exercises its intelligent capabilities within each step but does not autonomously decide what to do next.

Autonomous Agent:
An agent that, beyond the basic agent capabilities, can autonomously plan tasks, reason step by step, and dynamically adjust its execution strategy. The key distinction from agentic workflows is that an autonomous agent can independently determine its execution steps and strategy adjustments based on the task objective, rather than following a predefined pipeline. Its "autonomy" lies in its ability to dynamically formulate strategies and steps to complete a task.

Glossary

Interleaved Thinking:
A working mode in which a large language model interleaves reasoning, reflection, and planning during the process of generating output. The model can pause generation to analyze and assess available information, then decide what to output next, rather than generating all results in one continuous pass.

Skill/Tool:
A functional module that an agent can invoke during task execution, including capabilities such as legal database retrieval, document formatting, date calculation, and information cross-referencing. Skills and tools interface with the agent through standardized interfaces and are selected and invoked by the agent as needed. Where no ambiguity arises, "skills and tools" may be referred to simply as 'skills'.

Capability Component:
A general term for all types of capability units that an agent can invoke during task execution, encompassing skills and tools, plugins, external interfaces, and subagents. Skills and tools are the primary form of capability components; subagents are capability components that take the form of agents. The introduction, certification, and monitoring of capability components are key aspects of agent security management.

Subagent:
An agent instance created or invoked by a primary agent during task execution to complete a specific subtask. A subagent possesses the general capabilities of an agent and has its own task comprehension and execution process. Its behavioral boundaries and permissions should be constrained by the scope of the primary agent's task. This is the key distinction from skills and tools: skills and tools are functional modules that are invoked and do not have a decision-making process of their own. In multi-agent collaborative scenarios, anomalous behavior by a subagent can propagate and amplify along the task chain.

Harness Engineering:
An engineering methodology for building constraints, guardrails, feedback loops, and lifecycle management tools around an agent, with the goal of enabling the agent to consistently produce correct, auditable, and maintainable results in production environments. The core formula of harness engineering is "Agent = Model + Harness": the model provides reasoning capabilities, while the harness provides tool coordination, safety guardrails, error recovery, monitoring and transparency, and human intervention mechanisms. In judicial settings, harness engineering particularly emphasizes using judicial professional methodology to guide how AI technologies are combined and applied.

Glossary

Guardrail:
Security control measures for large language models and agents that use rule libraries, detection models, and other technical means to detect, intercept, and handle issues such as problematic input and output content, data leakage, and prompt attacks. Guardrails restrict agents from producing content that does not meet expectations, ensuring the controllability, compliance, and security of generated content. Guardrails are a core component of harness engineering and are typically deployed in multiple layers.

Sandbox:
A controlled, isolated environment provided for agent task execution. Sandboxes confine an agent's operations within authorized boundaries, preventing it from accessing unauthorized system resources or causing unintended effects on the external environment. The division of responsibility between sandboxes and guardrails is as follows: sandboxes isolate the scope of an agent's actions at the system level, while guardrails filter an agent's inputs and outputs at the content level.

Task:
A work instruction unit issued by a user to an agent, with a clear objective and expected output. A task typically includes elements such as the task objective, scope, skill requirements, and expected output. Tasks are classified into two categories by user involvement: predefined tasks, whose execution steps and invoked skills are preset by the system developer; and custom instruction tasks, where the user describes the requirements and the agent autonomously plans the execution path.

Context:
All information that an agent can access and reference during task execution, including case materials provided by the user, task instructions, system-preset knowledge, and intermediate results generated during task execution.

Glossary

Instruction/Prompt:
The text in which a user describes task requirements to an agent in natural language. It is the primary basis for the agent to understand the task objective and plan the execution path. In technical documentation, instructions are also referred to as prompts. The completeness and precision of an instruction directly affect the quality of the agent's execution.

Jailbreak:
An attack technique that uses role-playing, encoding bypasses, multi-turn gradual inducement, and other specific input methods to trick a large language model or agent into ignoring its built-in safety constraints and performing actions it should refuse. Jailbreaking is a form of prompt injection attack. In judicial settings, jailbreak attacks may be carried out through channels such as submitted litigation materials, leading questions, or customized parameter settings. Jailbreaks are a primary target for guardrail detection and interception.

Memory:
The mechanism by which an agent stores, manages, and retrieves information, enabling the agent to retain and utilize historical interaction data, task experience, and domain knowledge. Memory is divided into short-term memory (immediate information within the current task) and long-term memory (experience and knowledge accumulated across tasks). In judicial settings, memory management is especially critical: sharing memory across cases can lead to information cross-contamination and must be isolated by case and cleared by task.

Loop:
The working mode in which an agent repeatedly cycles through "perceive-reason-act-feedback" during task execution. Loops exist on two levels: the inner loop is the iterative cycle within a single task (the agent repeatedly invokes tools and refines intermediate results until the task objective is met); the outer loop is the experience-accumulation cycle across tasks (the user progressively optimizes instructions and supervision strategies through repeated use).

Appendix 3: Chronology of the Development of theShenzhen Court's AI-Assisted Adjudication System (as of August 2026)

2023

May.

Planning and construction of the AI-Assisted Adjudication System were initiated.

The four development principles — firmly position AI as a supporting tool; deeply integrate AI with the adjudication process; fully implement judicial accountability; thoroughly ensure data security — were established, along with the "build, use, and concurrently refine" approach.

Aug.

The Shenzhen Intermediate People's Court formally established the research and development mechanism.

Shenzhen Courts set up a leading group and a project office for AI-assisted adjudication applications, adopting a flat team structure combining full-time and part-time members. Work began on defining application rules, operational mechanisms, and requirement designs.

Nov.

The technical team fully moved on-site and formally commenced development.

Fifteen development groups and eleven primary court application groups were successively formed to carry out deep, integrated work with the technical team.

Dec.

President Zhang Jun of the Supreme People's Court (SPC) heard the first briefing and issued instructions.

President Zhang Jun put forward the work requirement of "not lying flat and daring to break through." The first development phase was concluded, and some functional modules were connected to the test environment.

2024

Jan.

The Intelligent Case-Filing Review Module was first launched for pilot operation.

The module covers 473 categories of civil and commercial case types, assisting judges in verifying all 35 filing standards item by item.

Apr.

Post-training of the first judicial domain-specific LLM was completed.

Based on a specialist legal corpus approaching a trillion Chinese characters (including knowledge graphs and data engineering from the databases) provided by the People's Court Press, and a judicial practice evaluation system developed by Shenzhen Courts, the post-training of the first judicial domain-specific LLM was completed with the pre-launch testing being passed.

May. 11

Shenzhen's "AI + Data×" On-Site Promotion Meeting was held in the Shenzhen Intermediate Court.

Shenzhen Intermediate People's Court hosted the on-site promotion meeting for the city's "AI + Data ×" initiative. The event was attended by leading municipal officials. During the event, the pilot version of the System, currently under internal testing, was demonstrated and presented to major municipal government departments.

Jun. 28

★Version 1.0 of the System was officially launched.

The System covered major civil case types, accounting for 90 percent of total adjudication caseloads. Meanwhile, the first judicial domain-specific LLM tailored for adjudication scenarios was officially activated, enabling AI to empower the entire adjudication process. This landmark achievement garnered extensive coverage from the media, including Xinhua News Agency, CCTV.com, and Nanfang Daily.

Oct. 15

★President of Shenzhen Intermediate People's Court delivered a special keynote speech in Singapore.

President Zhao Juhua of the Shenzhen Intermediate People's Court delivered a special keynote speech at the Eighth China-Singapore Legal and Judicial Roundtable, sharing the practices of judicial AI application. The Honourable the Chief Justice of Singapore expressed his hope to send a delegation to Shenzhen for a study visit in 2025.

Dec. 31

★Version 2.0 of the System was officially launched, achieving full coverage across civil, administrative, and criminal cases.

It expanded coverage from civil cases to all civil, administrative, and criminal cases. In the six months since launch, the System had assisted with approximately 152,000 case-file reviews, 63,000 hearings, and 145,000 judicial documents. The average case-closure time in Shenzhen Courts had been shortened by 38 days year-on-year, and the appeal rate had dropped by 24% year-on-year.

2025

Jan. 1

Xinhua News Agency released an English feature article to the world.

Xinhua News Agency published an English feature article, "China's Local Judicial Systems Embrace AI to Improve Efficiency," specifically highlighting the Shenzhen Courts' AI-Assisted Adjudication System.

Jan. 2

Harvard's Journal of Legal Analysis published a dedicated research paper.

The Journal of Legal Analysis , an open-access journal of Harvard Law School published by Oxford University Press, published a 25-page article titled "How do judges use large language models? Evidence from Shenzhen", which found that the Shenzhen court "... making it the first court in the world to routinely and extensively use generative AI to make judgments."

Apr. 23

The project was selected as one of Shenzhen's "Top 10 Rule of Law Events" for 2024.

In the annual "Rule of Law Events" campaign hosted by the Political and Legal Affairs Commission, the development of the System was selected as one of the "Top 10 Rule of Law Events" of 2024 in Shenzhen.

Aug.

A partnership with Huawei Technologies Co., Ltd. was established to create Huawei's global computing model site for AI-powered judicial applications.

The cooperation period spans from August 2025 to August 2027. Huawei has deployed a dedicated team to advance the engineering development of computing terminals and technical improvement of the Ascend ecosystem.

Sep. 16

★The National Court Judicial AI Application Joint Innovation (Shenzhen) Base was officially launched.

The SPC held a symposium on promoting the research and application of intelligent judicial technologies at the Shenzhen Intermediate People's Court. The National Court Judicial AI Application Joint Innovation (Shenzhen) Base was officially launched, marking the nation's sole base of its kind. The base is tasked with serving as both a "laboratory" and a "testing ground" for intelligent applications in the judiciary.

Nov. 2

The "Pengfa Tianping Cup" AI-Assisted Adjudication Skills Competition finals were held.

Shenzhen Courts partnered with the Shenzhen Federation of Trade Unions and the Human Resources and Social Security Bureau of Shenzhen Municipality to organize a citywide competition in AI-assisted adjudication. More than 4,100 judges and other court personnel took part—the largest turnout for a skills competition in Shenzhen.

Nov. 17

Chief Justice of Singapore led a delegation to visit Shenzhen Courts.

After observing a live demonstration of the System, the Chief Justice of Singapore remarked: "This is the most impressive demonstration I've seen in any jurisdiction. The system is very advanced and sophisticated... This is where the world is moving forward."

2026

Jan. 22

★Phase II of the AI-assisted adjudication project officially commenced, introducing LLM-powered autonomous agents to support the iteration.

The second phase of the AI-assisted adjudication project was officially launched, adopting LLM-powered autonomous agents to support the iteration and development of upgraded system versions. Taking labor dispute cases as the pilot type , the project has established an "AI agent + skill tool" model, with a layout to cover 90 percent of common case types in the future.

Mar. 10

★President Zhang Jun of the SPC affirmed the "Shenzhen experience" during a National People's Congress session.

During the deliberation of the SPC work report at the Guangdong Delegation panel meeting of the Fourth Session of the 14th National People's Congress, President Zhang Jun affirmed Shenzhen's AI-Assisted Adjudication System. He noted that the SPC had launched pilot operations at 23 intermediate and primary courts in 11 provinces (autonomous regions and municipalities), with "excellent results."

Apr. 21

The world's first Guide for Using AI Agents in the Judicial Domain was released in Shenzhen.

Jointly compiled by the Shenzhen Intermediate People's Court, Shenzhen Research Institute of Big Data (SRIBD), the Faculty of Law, The University of Hong Kong and Institute for Internet Judiciary, Tsinghua University (THUIJ). The Guide clarified the four principles and the four negative lists. On the same day, SRIBD and the Shenzhen Intermediate People's Court signed a strategic cooperation agreement.

Apr. 24

★The Labor Cases Tailored Agent, the General-purpose Assisting Agent, and the Case-handling Assisting Agent were launched.

The Labor Cases Tailored Agent was officially upgraded and launched, featuring functions such as case summary generation, delivery of key adjudication points, and guided document drafting. The General-purpose Assisting Agent and the Case-handling Assisting Agent for all civil and commercial case types were simultaneously launched. Judges can directly input instructions into the AI agents, which can automatically identify and invoke corresponding expertise skills to assist case handling.

Apr. 27

A keynote speech titled "From Shenzhen to the Future" was delivered at the Guangdong Province AI Application Matchmaking Conference.

Deputy President Kuang Xiaohua of the Shenzhen Intermediate People's Court delivered the speech "From Shenzhen to the Future" at the Guangdong Provincial AI Application Matchmaking Conference, highlighting that in just over a year since its launch, the System had deeply supported the adjudication of more than 600,000 cases.

May. 16

The Case-handling Assisting Agent was extended to all case types, covering civil, administrative, and criminal proceedings.

The General-purpose Assisting Agent was expanded to cover all case types (including administrative and criminal). The Case-handling Assisting Agent has been iterated and upgraded with new functions, including knowledge query, legal research, and knowledge skills access.

Jun. 5

The System obtained an agent security capability assessment certificate.

The System's Agent passed the Agent Security Capability Assessment conducted by China Telecommunication Technology Labs (CTTL) under the China Academy of Information and Communications Technology (CAICT), achieving a Level 5 (highest-level) security rating.

Jun. 13

★Administrative and criminal Case-handling Assisting Agents were launched, achieving full case-type coverage.

The administrative and criminal Case-handling Assisting Agents were officially launched, meaning Case-handling Assisting Agents now cover all case types across the three types of adjudication cases.

Jul. 24

★"Specialized Skill Packs" for seven major civil and commercial case types were officially launched.

The packs cover first- and second-instance cases for seven case types, including private lending disputes. Five core skills were launched: evidence summarization and evaluation, generation of findings of fact, extraction and computation of calculation information, matching of disputed issues to key adjudication points, and reasoning style replication. Judges can input direct instructions to the AI agent, and the Agent will automatically identify and invoke the corresponding expertise skills, or obtain refined outputs via the "Skills" button.