

Overview of the 1st Workshop on NLP for Empowering Justice

Ashutosh Modi¹, Saptarshi Ghosh², Asif Ekbal³, Pawan Goyal², Sarika Jain⁴,
Abhinav Joshi¹, Shivani Mishra¹, Debtanu Datta², Shounak Paul²,
Kshetrimayum Boynao Singh³, Sandeep Kumar³

¹IIT Kanpur, India, ²IIT Kharagpur, India, ³IIT Patna, India, ⁴NIT Kurukshetra, India

Correspondence: ashutoshm@cse.iitk.ac.in

Abstract

The first iteration of the JUST-NLP: Workshop on NLP for Empowering Justice was organized to accelerate research in Natural Language Processing for legal text processing. The inaugural edition, JUST-NLP 2025, was held as a hybrid event at IJCNLP-AAACL 2025 on December 24 at IIT Bombay. The program featured a research track, four invited talks, and two shared tasks: (1) L-SUMM, an abstractive summarization task for Indian legal judgments, and (2) L-MT, a legal machine translation task between English and Hindi. The workshop received strong interest from the community, with 29 submissions, of which 21 were accepted. Among the accepted papers, 5 were regular research-track papers published in the proceedings, and 2 were accepted as non-archival presentations. For the shared tasks, 9 papers were accepted for L-SUMM, and 5 papers were accepted for L-MT, for publication in the proceedings. The workshop focused on a broad set of Legal NLP challenges, including information extraction, retrieval, multilingual processing, legal reasoning, and applications of large language models. Overall, JUST-NLP 2025 aimed to bring together AI researchers and legal practitioners to develop scalable, domain-aware NLP methods that can support legal workflows and contribute toward more efficient and equitable justice systems.

1 Introduction

Legal documents form the textual backbone through which societies articulate rights, enforce responsibilities, and administer justice. They encode statutory frameworks, capture the reasoning behind judicial decisions, and document complex interactions among individuals, institutions, and the state. As a result, Legal NLP has emerged in recent years as a significant research area: legal texts are not only primary to ensuring fairness, transparency, and accountability in democratic systems, but also indispensable for facilitating faster

and more structured judicial processes. However, these documents pose challenges that are far more demanding than those encountered in many other NLP domains, owing to their specialized language, intricate reasoning structures, and strong reliance on precedent and contextual interpretation.

Legal documents often require a precise and unambiguous understanding, a high level of factual and logical consistency, and substantial background/contextual knowledge for processing. Moreover, these documents span hundreds of pages, follow rigid structural conventions, and contain dense argumentation supported by layered citations to statutes, procedural rules, and prior cases. In multilingual jurisdictions such as India, these challenges are even greater, where legal information is spread across multiple languages, various court levels, and diverse document types, including judgments, petitions, orders, and statutes, each with its own linguistic and semantic characteristics. Such complexity motivates the need for domain-specific models, datasets, and evaluation protocols, as general-purpose NLP methods frequently fail to capture these nuances.

Countries with large populations, most notably India, with over 44 million pending cases, struggle with delays driven by manual workflows, difficulty in locating relevant precedents, and limited accessibility of legal information for both practitioners and citizens. Advances in NLP and large language models provide a promising avenue to mitigate these bottlenecks through tasks such as automated summarization, translation, precedent retrieval, and assistive legal reasoning. However, progress has been slowed by the scarcity of annotated legal datasets, the limited collaboration between technologists and legal experts, and the particularly high stakes of deploying AI in legal settings.

The JUST-NLP workshop was created to address these challenges by bringing together researchers from NLP, information retrieval, machine learn-

ing, and AI ethics, alongside law practitioners and scholars. Its goal is to provide a dedicated venue for the development of domain-aware models, high-quality legal datasets, multilingual resources, and discussions on the risks and opportunities of AI for justice. JUST-NLP invited research contributions spanning legal reasoning, information extraction, statute and precedent retrieval, multilingual legal processing, and applications of large language models in legal workflows. Overall, the program featured a research track, four invited talks (detailed in Section 3), and two shared tasks: L-SUMM and L-MT (detailed in Section 4).

As the first edition of the workshop co-located with IJCNLP-AAACL 2025, JUST-NLP aims to establish a sustained, interdisciplinary forum that promotes both the scientific advancement of Legal NLP and its responsible deployment toward more efficient, inclusive, and equitable justice systems.

2 Program

The first iteration of the JUST-NLP 2025 workshop featured a research track with open submissions, four invited talks, and two shared task tracks. The workshop received 29 submissions, of which 21 were accepted, highlighting a strong community interest in Legal NLP and its applications to the legal processing pipeline. Among the accepted papers, 5 were regular research-track papers published in the proceedings, and 2 were accepted as non-archival presentations. For the shared tasks, 9 papers were accepted for L-SUMM, and 5 papers were accepted for L-MT, for publication in the proceedings. Since the event was held in a hybrid format, presentations were delivered both in person at the IJCNLP-AAACL 2025 venue in Mumbai and virtually.

The accepted papers provided a broad spectrum of Legal NLP research, highlighting community interest in both foundational challenges and emerging directions in the field of legal text processing. The contributions/submissions aim to address legal information extraction, knowledge graph construction, statute and precedent retrieval, legal citation network modeling, consumer-law assistance systems, and specialized pipelines for legal analytics. A significant portion of the accepted work focused on multilingual and English–Hindi legal machine translation, as well as a wide range of approaches to legal document summarization, including hierarchical and structure-aware chunking, retrieval-augmented, and agentic LLM work-

flows. Across these efforts, a clear trend emerged as the integration of large language models into legal tasks, with a particular emphasis on long-document processing, domain adaptation, reinforcement learning–based fine–tuning, and hybrid extractive–generative strategies. Overall, these papers highlight the community’s growing interest in scalable, transparent, and domain-grounded LLM methods for real-world legal reasoning and decision support.

In addition to the research track, the workshop also hosted two shared tasks, L-SUMM, an abstractive summarization task for Indian legal judgments, and L-MT, a legal machine translation task between English and Hindi. Both tasks demonstrated active participation from multiple research groups, highlighting the growing interest in this area.

3 Invited Talks

The JUST-NLP 2025 workshop featured a diverse set of invited talks from leading experts across law, industry, public institutions, and AI research. The talks highlighted multiple perspectives on the future of Legal NLP, ranging from legal theory and courtroom-scale AI deployment to ontology-driven knowledge representations and technology-enabled governance reform.

3.1 Prof. Niraj Kumar

Affiliation: National Law University Delhi

Title: Artificiality of Law in Reference to Artificial Intelligence

Summary: Prof. Niraj Kumar delivered a keynote examining the evolving relationship between legal theory and emerging AI systems. Drawing on his expertise in constitutional and comparative law, he reflected on how AI challenges traditional conceptions of legal reasoning, authority, and legitimacy. His talk explored foundational questions around the “artificiality” of law itself and the implications of entrusting interpretive or decision-support roles to artificial agents.

3.2 Mr. Arghya Bhattacharya

Affiliation: Co-Founder & CTO, Adalat AI

Title: Building AI for India’s Courtrooms: Scaling Sovereign Voice & Language Systems Across 20% of the Judiciary

Summary: Mr. Arghya Bhattacharya shared the technical and operational journey behind deploying large-scale voice and language AI across thousands of Indian courtrooms. His talk detailed the

engineering of sovereign, privacy-preserving systems for live transcription, dictation, translation, and legal assistance, addressing challenges such as courtroom acoustics, multilingual phrasing, domain grounding, and reliable scaling in high-stakes public institutions. The session also outlined how voice AI forms part of a broader “AI for Justice” ecosystem and discussed pathways for building inclusive digital public goods for the Global South.

3.3 Mr. Joseph Pookkatt, Ms. Sampritha Manjunath, and Mr. Parth Parikh

Affiliation: Staram Analytics & eSuccess AI Technologies

Title: Designing and Implementing Knowledge Graphs in the Legal Domain

Summary: This invited industry panel examined the role of ontologies and knowledge graphs in enabling robust, interpretable, and domain-grounded AI systems for legal applications. The speakers highlighted why traditional RAG pipelines struggle with semantically distant but logically related concepts and demonstrated how GraphRAG and structured symbolic layers can significantly improve retrieval and reasoning. Using Indian legal use cases, they showcased approaches for building smaller, specialized legal knowledge graphs and discussed how LLMs can assist in graph architecture inference, supporting next-generation legal research, compliance, and decision-making systems.

3.4 Ms. Nishi Yadav

Affiliation: Senior Legal Consultant, Ministry of Tribal Affairs, Government of India

Title: Fixing the Process, Powering the System: What Jharkhand Teaches India About Legal-Tech

Summary: Ms. Nishi Yadav presented a governance-focused perspective on legal-tech transformation, using Jharkhand’s pioneering reforms as a case study. Her talk argued that sustainable legal-tech begins with process redesign, mapping institutional workflows, fixing structural bottlenecks, and only then introducing technology. The session traced Jharkhand’s evolution from the Vidhi Portal to the Integrated Litigation Management System (ILMS), highlighting sharp reductions in case pendency, faster response cycles, and cultural shifts toward data-driven governance. She also outlined the state’s next phase: integrating AI-enabled predictive governance to identify systemic drivers of litigation and strengthen policy design.

4 Shared Tasks

JUST-NLP 2025 hosted two shared tasks, each organized by members of the workshop organizing committee. Detailed descriptions, datasets, baselines, and system analyses are provided in the individual shared task overview papers (Datta et al., 2025; Singh et al., 2025) included in the proceedings.

Legal Summarization (L-SUMM) The L-SUMM (Datta et al., 2025) task focused on abstractive summarization of Indian legal judgments. Participants were required to generate concise and coherent summaries that capture the core legal reasoning and outcomes of complex, lengthy judicial documents. The task highlighted key challenges in legal summarization, including domain-specific terminology, multi-paragraph argumentation structures, and the need for faithful condensation of legal rationale. Evaluation was conducted using ROUGE-2, ROUGE-L, and BLEU, and participating teams explored a variety of approaches, including domain-adapted LLMs and long-context transformers. Full details and results appear in the corresponding task overview paper (Datta et al., 2025).

Legal Machine Translation (L-MT) The L-MT (Singh et al., 2025) task targeted English–Hindi legal machine translation, addressing the need for bilingual accessibility in India’s multilingual judicial system. The task required systems to handle complex legal syntax and terminology while preserving semantic and legal fidelity. Submissions were evaluated using BLEU, METEOR, TER, chrF++, BERTScore, and COMET. Approaches ranged from fine-tuned encoder–decoder models to instruction-tuned large language models. A separate overview paper (Singh et al., 2025) in this volume provides a detailed presentation of the dataset, baselines, and performance analysis.

5 Workshop Overview and Outlook

The organizers were encouraged by the strong and diverse response to the inaugural JUST-NLP workshop. Accepted contributions spanned a wide spectrum, from applied systems addressing legal document summarization, translation, and knowledge retrieval, to foundational research exploring legal reasoning, multilingual NLP, and domain-adapted large language models. This breadth demonstrates both the societal relevance of NLP for justice and the technical richness of the domain.

JUST-NLP has helped bring together an emerg-

ing research community focused on legal NLP, broad enough to foster interdisciplinary collaboration, yet focused enough to make rapid progress on high-impact challenges. The workshop sits at the intersection of NLP, information retrieval, AI for governance, and legal studies, reflecting the need for cross-disciplinary approaches to improve access to justice.

Looking ahead, the workshop aims to continue expanding both its technical and thematic scope. Future editions will encourage contributions on fairness, explainability, and ethical AI in legal systems, as well as multilingual, cross-jurisdictional, and low-resource challenges. Building on the success of the shared tasks in legal summarization and machine translation, future iterations will continue to provide concrete problem settings that engage the research community and address pressing societal needs. The organizers hope that JUST-NLP will grow into a sustained forum for innovation, collaboration, and responsible deployment of NLP technologies in support of equitable and efficient justice systems.

6 Conclusion

Legal documents present a rich and challenging domain for NLP research, with high societal impact and pressing real-world applications. While prior work in information retrieval, knowledge representation, and computational linguistics provides a foundation, many critical challenges remain, ranging from understanding complex reasoning in judgments to handling multilingual and low-resource legal contexts. The automated processing of legal texts is still in its early stages, and careful, domain-aware approaches are essential to ensure accuracy, fairness, and reliability.

By providing a dedicated forum for interdisciplinary collaboration, the inaugural JUST-NLP workshop has highlighted these challenges and offered opportunities for researchers and practitioners to exchange ideas, benchmark systems, and explore innovative solutions. We hope that future editions of JUST-NLP will continue to foster the development of domain-specific models, curated datasets, and practical applications, ultimately contributing to a more efficient, accessible, and equitable justice system in India and beyond.

References

- Debtanu Datta, Shounak Paul, Kshetrimayum Boynao Singh, Sandeep Kumar, Abhinav Joshi, Shivani Mishra, Sarika Jain, Asif Ekbal, Pawan Goyal, Ashutosh Modi, and Saptarshi Ghosh. 2025. Findings of the JUST-NLP 2025 Shared Task on Summarization of Indian Court Judgments. In *Proceedings of the First Workshop on NLP for Empowering Justice (JUST-NLP)*, Mumbai, India. Association for Computational Linguistics.
- Kshetrimayum Boynao Singh, Sandeep Kumar, Debtanu Datta, Abhinav Joshi, Shivani Mishra, Shounak Paul, Pawan Goyal, Sarika Jain, Saptarshi Ghosh, Ashutosh Modi, and Asif Ekbal. 2025. Findings of the JUST-NLP 2025: A Shared Task on English-to-Hindi Legal Machine Translation. In *Proceedings of the First Workshop on NLP for Empowering Justice (JUST-NLP)*, Mumbai, India. Association for Computational Linguistics.