NyayaSakhi–SWATI: India's First Statute-Aligned, Retrieval-Augmented LAMP² 4.0 AI-Powered Digital Legal Companion for Victims of Domestic-Violence

Dipali Awasekar [1], Dr.L.M.R.J.Lobo[2]
[1]Information Technology Dept., Walchand Institute of Technology, Solapur, Maharashtra
[2] Information Technology Dept., Walchand Institute of Technology, Solapur, Maharashtra
[1] dipali.awasekar@gmail.com
[2] headitwit@gmail.com

***Abstract*—Domestic violence (DV) remains a pervasive challenge in India, and survivors frequently struggle to obtain timely, comprehensible, and statute-accurate information about their rights under the Protection of Women from Domestic Violence Act, 2005 (PWDVA). Existing helplines and portals are typically static, FAQ-based, or rule-driven, offering neither interactive legal literacy nor case-specific guidance grounded in statutory relief structures. This paper presents NyayaSakhi–SWATI, India's first statute-aligned, retrieval-augmented AI legal companion for DV survivors, powered by the LAMP² 4.0 backend. SWATI (Support and Welfare Assistance through Technology Interface) is designed as an empathetic "AI NyayaSakhi" that securely collects survivor narratives, retrieves PWDVA-aligned precedents from the NyayaSmriti vector knowledge base, and delivers plain-language responses on likely statutory reliefs, indicative case duration, and prescriptive next steps. At the core of the system is the Statute-Aligned Legal Relief Prediction (SALRP) task, wherein LAMP² 4.0 predicts multi-relief outcomes (protection, residence, custody, monetary relief, and compensation) directly from narrative inputs. On the NyayaSmriti evaluation set, LAMP² 4.0 achieves an accuracy of 81%, outperforming strong Indian legal-AI baselines including INLegalLlama, PredEx, and NyayaRAG by margins of approximately +8–9 percentage points and +13.93 percentage points respectively. A single-case reference evaluation against an adjudicated PWDVA matter demonstrates complete alignment between predicted and judicially granted reliefs. A research-preview study with domestic-violence survivors and experienced PWDVA practitioners further indicates that NyayaSakhi–SWATI is perceived as usable, empathetic, and statute-faithful. Together, these findings show how retrieval-augmented, statute-aligned AI can move beyond judgment-centric prediction toward survivor-centric, pre-adjudicative decision support in domestic-violence litigation.**

***Keywords*—Artificial intelligence; Domestic violence; Retrieval-augmented generation; Statutory relief prediction**

ICTIEE* Track—*Emerging Technologies and Future Skills
ICTIEE Sub-Track— AI, Machine Learning, and Digital Tools in Education

I. INTRODUCTION

DOMESTIC violence (DV) remains one of the most pervasive human-rights violations in India, with severe consequences for the physical, psychological, and economic well-being of women (Ganguli, 2024). The Protection of Women from Domestic Violence Act, 2005 (PWDVA), was enacted to provide a civil, survivor-oriented framework of remedies, including protection orders, residence rights, custody arrangements, monetary relief, and compensation (Government of India, 2005). However, despite the progressive architecture of the law, access to these statutory entitlements remains deeply uneven. Many survivors face a persistent justice gap arising from low awareness of rights, entrenched social stigma, procedural complexity, and limited access to legal-aid services (Bankata & Mishra, 2025).

This access gap is compounded by a heavily burdened judicial system. As of 2025, the National Judicial Data Grid reports that over 4.73 crore cases are pending in India's district courts. Of these, 3.62 crore (approximately 76%) are criminal matters and 1.10 crore (24%) are civil. Critically, more than 1.83 crore of these cases—around 8% of the total—were filed by women, many involving family disputes and domestic violence. Furthermore, over 62.8% of criminal and 58.5% of civil cases have been pending for more than a year, delaying access to justice and undermining timely relief, particularly for survivors seeking urgent protection under the PWDVA (National Judicial Data Grid, 2025).

In this context, digital technologies have increasingly been explored as instruments to promote civic awareness, legal literacy, and access to justice. Among these, conversational AI systems, or chatbots, offer unique promise. Their ability to engage users in personalized, interactive dialogue allows for the delivery of information in formats that reduce cognitive load and accommodate varying literacy levels. Prior work has demonstrated that dialogue-based systems can enhance learning, reduce informational asymmetries, and help users navigate complex institutional processes by simulating natural conversations and providing stepwise guidance (Awasekar & Lobo, 2024a). In sectors such as health and education, large-scale chatbot deployments have proven effective in reaching marginalized groups with multilingual, persona-based, and mobile-friendly interfaces (Awasekar & Lobo, 2025). However, in the Indian legal domain—and specifically in the context of domestic violence—chatbot technology remains largely underutilized, with existing tools limited to static FAQs, generic grievance mechanisms, or non-statutory advice

To address this critical gap, we introduce NyayaSakhi–SWATI, a digital legal companion designed specifically to support women facing domestic violence. The system integrates conversational AI with statute-aligned reasoning to deliver legal literacy and actionable guidance focused on the PWDVA. It is instantiated through "Swati," an empathetic female persona that engages in multi-turn conversations to elicit key case facts and return legally grounded, accessible advice. Unlike static awareness platforms, NyayaSakhi–SWATI dynamically processes user narratives and transforms them into structured legal representations and predicted relief outcomes under PWDVA Sections 18–22.

The technical foundation of NyayaSakhi–SWATI lies in the LAMP² 4.0 framework, which constitutes a Retrieval-Augmented Generation (RAG) architecture at the core of a broader AI pipeline for Statute-Aligned Legal Relief Prediction (SALRP). This framework incorporates curated legal corporaNyayaDeepa, NyayaPariksha, and NyayaSmriti alongside a legal-tuned LLaMA 3.1 model and vector database, enabling the retrieval of relevant statutory provisions and case precedents and the generation of statute-aligned advisory content (Nigam et al., 2024; Nigam et al., 2025a; Nigam et al., 2025b). This architecture allows for narrative-level, statute-grounded predictions of multi-relief outcomes, closing the semantic gap between survivor experiences and legal provisions. Importantly, the SALRP formulation diverges from traditional Legal Judgment Prediction (LJP), which is focused on appellate outcomes; instead, SALRP emphasizes pre-adjudicative guidance based on first-instance complaint narratives and trial-level judgments.

This paper presents the design, deployment, and empirical evaluation of NyayaSakhi–SWATI along three research dimensions. First, it assesses whether the LAMP² 4.0 Retrieval-Augmented architecture embedded within SWATI outperforms existing Indian legal-AI baselines in statute-aligned relief prediction for DV cases (RQ1). Second, it evaluates the system's ability to accurately replicate relief outcomes in previously adjudicated PWDVA cases with known judicial judgments (RQ2). Third, it examines how survivors and legal practitioners perceive the system's usability, trustworthiness, and helpfulness in a research-preview setting (RQ3).

By addressing these questions, NyayaSakhi–SWATI expands the frontier of legal AI beyond traditional judgment prediction, offering a novel paradigm of survivor-centric, statute-aligned digital legal support. It demonstrates that retrieval-augmented, domain-tuned language models can be combined with empathetic conversational agents to provide legally accurate, safe, and explainable assistance for women navigating domestic violence law in India.

II. RELATED WORK

Recent developments in legal AI have produced a range of systems aimed at automating legal judgment prediction, summarization, and information retrieval. In India, early efforts focused on classical machine learning models—such as SVMs and decision trees—for outcome prediction using appellate court data (Malik et al., 2021). While these models achieved high accuracy in binary classification tasks, they depend heavily on structured metadata and are inherently unsuitable for tasks requiring statute-aligned reasoning, such as mapping user narratives to specific sections of the Protection of Women from Domestic Violence Act, 2005 (PWDVA) (Government of India, 2005). More recently, transformer-based approaches have improved the parsing of judicial texts. Legal case summarization models (Shukla et al., 2022) and large-scale datasets like NYAYAANUMANA (Nigam et al., 2025) and ILDC-CJPE (Malik et al., 2021) have significantly contributed to judgment prediction and explanation. However, these systems remain post-adjudicative, trained on appellate rulings, and are not designed for interpreting unstructured, first-person complaint narratives or for pre-adjudicative guidance.

RAG-based systems represent a more recent evolution in legal AI, combining neural language models with retrieval mechanisms to enhance legal recall and factual grounding. Projects such as JusticeNetBD (Bangladesh), and NyayaRAG (India) exemplify how integrating curated corpora with neural retrievers enables legal explainability and conversational assistance (Nigam et al., 2025). Despite their promise, these tools are typically built for retrospective analysis and rarely align with civil remedies under statutes like the PWDVA. Moreover, adversarial vulnerabilities in the retrieval layer of RAG pipelines pose additional challenges in high-stakes legal settings. Similarly, corpus-building initiatives like ILDC, ICLC, and ECHR have advanced NLP research in law but concentrate predominantly on High Court and Supreme Court decisions. These datasets often lack section-level tagging (e.g., §§18–22 of the PWDVA) and are not optimized for district or family court use, where domestic violence cases are typically adjudicated (Bankata & Mishra, 2025).

The gap in legal-AI systems that cater to vulnerable populations is further evident in the design of public-facing legal chatbots. Tools such as OpenNyAI, LawRato, and Sama provide static FAQ interfaces or directory-based legal help but lack predictive capabilities or statute-aware reasoning. They are not designed to interpret complex narratives or offer multi-relief estimations. Conversely, persona-based chatbots like SnehAI have successfully delivered empathetic health education demonstrating the importance of trust-building interfaces for sensitive topics. However, these models do not operate within legal domains or connect users to statutory provisions. Government initiatives such as MyGov and UMANG prioritize grievance redressal and administrative workflows but are not designed for statute-aligned legal inference, especially for survivors of domestic violence

(National Judicial Data Grid, 2025).

This landscape reveals a critical and under-addressed set of

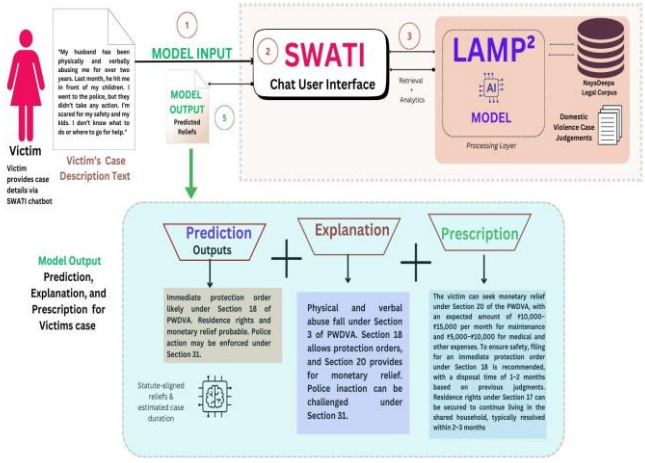

Fig. 1. End-to-end LAMP²–SWATI pipeline integrating user input, RAG retrieval, and legal advisory generation.

gaps in existing literature and deployed systems. First, there is an absence of legal-AI tools capable of multi-label, statute-aligned relief prediction—i.e., translating user narratives into specific legal entitlements under PWDVA Sections 18–22. Second, available datasets are overwhelmingly post-adjudicative, with sparse representation of the lower judiciary, making them unsuitable for early-stage legal prediction or survivor guidance. Third, there is a complete lack of empathetic, trauma-informed chatbot design that integrates RAG-based legal reasoning tailored to the needs of domestic violence survivors. Finally, existing systems are advisory or administrative in scope, rather than predictive, statutory, or narrative-aligned, and have not been validated through survivor interaction or legal clinics in real-world conditions.

To directly address these deficiencies, NyayaSakhi–SWATI emerges as India's first AI-based legal companion explicitly designed for survivors of domestic violence. Built on the LAMP² 4.0 Retrieval-Augmented Generation (RAG) framework, it leverages curated legal corpora—including NyayaDeepa, NyayaPariksha, and NyayaSmriti—to perform statute-aligned, pre-adjudicative reasoning and predict possible reliefs under PWDVA. The chatbot is embodied through a female-presenting persona, "Swati", designed to foster empathy, trust, and confidentiality—key for safe engagement with survivors. Unlike previous tools, NyayaSakhi transforms unstructured complaint narratives into structured legal representations and provides multi-relief guidance aligned with Sections 18–22. Confidential pilot deployments with survivor groups and legal aid clinics have demonstrated the system's usability, ethical design, and potential for real-world impact.

III. NyayaSakhi–SWATI: System Overview

NyayaSakhi–SWATI is conceived as a digital legal companion for women experiencing domestic violence, built to operationalise the LAMP² 4.0 phase of the Statute-Aligned Legal Relief Prediction (SALRP) framework. SWATI (Support and Welfare Assistance through Technology Interface) brings together three tightly coupled components: a retrieval-augmented legal reasoning engine aligned with the Protection of Women from Domestic Violence Act, 2005 (PWDVA), a statute-aligned gold-standard judgment corpus, and an empathetic conversational persona for survivor-facing interaction. The system is designed not as a generic chatbot, but as a statute-aware assistant that can listen to survivor narratives, translate them into legally meaningful features, and return grounded guidance on likely statutory reliefs and procedural next steps under PWDVA. As shown in Figure 1, survivors begin by entering a free-text description of their situation into the SWATI chat interface. The narrative is normalised and encoded into dense text embeddings using transformer-based encoders, and these embeddings are used to query NyayaSmriti, a gold-standard, statute-aligned subset of domestic-violence judgments curated for the SALRP task. The chatbot's legal knowledge base is derived from the

NyayaDeepa corpus, comprising 214,983 domestic-violence judgments collected from District and Family Courts across India, with major representation from Maharashtra. From this, an expert-annotated subset of 5,000 judgments, titled NyayaPariksha, was curated for statutory-relief labelling across Sections 18 (Protection Order) to 22 (Compensation Order) of the PWDVA. For the RAG phase, a gold-standard evaluation and retrieval set designated as NyayaSmriti was compiled as the statute-aligned SALRP benchmark. Each NyayaSmriti case is manually structured into a narrative case story and converted into dense vector embeddings within a Pinecone database, with summarisation using Gemini-1.5-Flash-8B to preserve legally salient details while maintaining retrieval efficiency and tight alignment with the PWDVA relief schema.

The retrieval layer returns statute sections and precedent fragments that are semantically close to the survivor's narrative. These are passed, along with the original narrative, to a legal-tuned LLM that implements the LAMP² 4.0 Retrieval-Augmented Generation (RAG) engine. This backend composes a statute-grounded response consisting of three elements: a bundle of predicted statutory reliefs (for example, protection, residence, monetary relief, custody, compensation), an accompanying explanation that links these reliefs to statutory provisions and similar precedents, and a prescriptive block

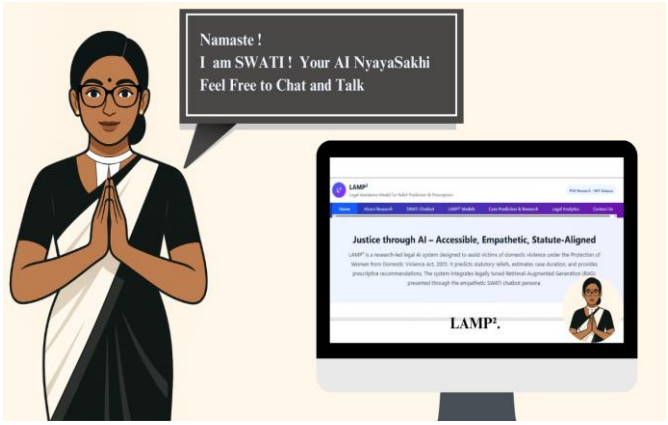

Fig. 2. SWATI conversational interface illustrating narrative intake and statute-aligned advisory output.

outlining indicative case duration, documentation checklists, filing steps, and safety-oriented suggestions.

The frontend is realized through the SWATI persona, illustrated in Figure 2. SWATI appears as a culturally grounded female avatar and mediates all interaction in a respectful, non-judgmental tone. The interface is implemented as a lightweight, web-based chat environment suitable for standard devices and typical connectivity conditions, with multilingual support (for example, English, Hindi, and Marathi) to reflect the linguistic realities of Indian survivors. Conversational flows are designed to be simple and legally accurate, with SWATI asking gentle clarifying questions, explaining relevant PWDVA provisions in accessible language, and summarising the model's predictions and prescriptions in a way that survivors can act upon. Privacy-by-design safeguards, including explicit consent prompts and anonymisation of conversational content in research-preview deployments, are integrated into the dialogue so that users can seek information without fear of exposure or judgment.

In the remainder of this paper, NyayaSakhi–SWATI is evaluated as a complete socio-technical system built on this architecture. The next section presents experimental results that examine: (i) how the LAMP² 4.0 backend compares with existing Indian legal-AI baselines on a statute-aligned benchmark; (ii) how its predictions align with a fully adjudicated PWDVA judgment; and (iii) how survivors and experienced PWDVA lawyers perceive the usability and helpfulness of the deployed SWATI interface in a research-preview setting.

IV. Experimental Evaluation and Results

This section presents the empirical evaluation of the LAMP² 4.0 RAG–Legal Tuned model as the core predictive engine for NyayaSakhi–SWATI. All experiments are conducted on the NyayaSmriti evaluation set, a gold-curated, statute-aligned benchmark of domestic-violence judgments derived from the ProtectDV corpus and specifically annotated for the Statute-Aligned Legal Relief Prediction (SALRP) task. Section 5.1 reports comparative benchmarking against state-of-practice Indian legal-AI models (addressing RQ1), Section 5.2 provides a fine-grained, single-case reference evaluation against a fully adjudicated PWDVA judgment (addressing RQ2), and Section 5.3 reports research-preview testing with survivors and expert lawyers (addressing RQ3).

A. RQ1 – Comparative Benchmarking of LAMP² 4.0 Against Indian Legal-AI Baselines

RQ1. To what extent does the LAMP² 4.0 Retrieval-Augmented Generation (RAG) architecture—operationalised through the NyayaSakhi–SWATI conversational legal companion—outperform existing Indian legal-AI models in delivering accurate, statute-aligned statutory relief predictions and empathetic advisory responses for domestic-violence survivors, while functioning as a fully deployed, real-world system rather than a laboratory-only prototype?

Before being embedded within the NyayaSakhi–SWATI conversational companion, the LAMP² 4.0 RAG–Legal Tuned architecture was rigorously benchmarked against state-of-practice Indian legal-AI systems. The objective of this evaluation was twofold: (i) to assess whether the proposed model offers a measurable advantage over existing models on a statute-aligned domestic-violence benchmark, and (ii) to justify its selection as the backend engine for a survivor-facing, AI-for-social-good deployment rather than retaining it as a purely experimental prototype. For this purpose, LAMP² 4.0 was compared with four representative baselines—Aalap (legal summarisation/translation), INLegalLlama (generic Indian legal LLM for legal Q&A and reasoning), PredEx (pre-RAG Legal Judgment Prediction of final verdicts), and NyayaRAG (RAG-based LJP for outcome plus explanation)—on the NyayaSmriti evaluation set, the first gold-curated, statute-aligned dataset of domestic-violence judgments annotated for multi-relief outcomes.

To situate LAMP²–SWATI within the broader Indian legal-AI landscape, the LAMP² 4.0 RAG–Legal Tuned model was benchmarked against the above baselines, each of which is considered strong within its respective task category. Aalap targets legal summarisation and translation; INLegalLlama is tuned for generic Indian legal reasoning (Q&A, bail reasoning, act/section understanding); PredEx focuses on pre-RAG Legal Judgment Prediction (final verdict: allowed/dismissed); and NyayaRAG is a RAG-based LJP system that outputs judgment labels with brief explanations. However, none of these models is explicitly designed for pre-adjudicative, statute-aligned relief prediction under PWDVA or for deployment within a survivor-centric conversational interface.

In contrast, LAMP² 4.0 is purpose-built for the SALRP task and evaluated on NyayaSmriti, which encodes multi-relief outcomes (protection, residence, monetary relief, custody, compensation, etc.) alongside narrative context. The consolidated metric-wise comparison (ROUGE, BLEU, BERT-F1, precision, recall, F1-score, accuracy) is visualised in Figure 3, and Figure 4, which together depict LAMP² 4.0's predictive and generative performance relative to the baselines. On the NyayaSmriti evaluation set, LAMP² 4.0 attains an accuracy of 81%, thereby establishing a new benchmark for early-stage statutory relief prediction.

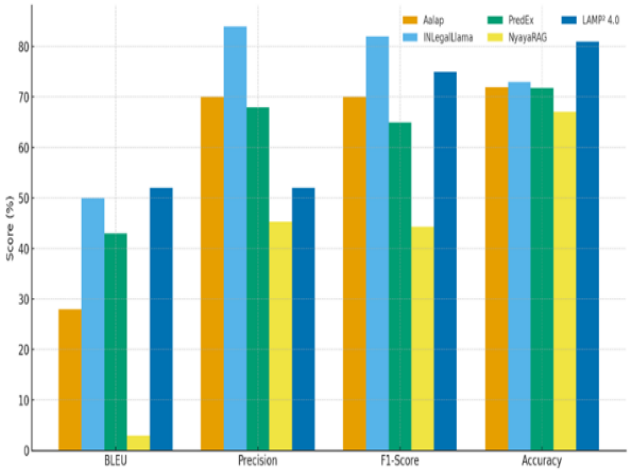

Fig. 3. ROUGE, BLEU, and BERT-F1 comparison of LAMP² 4.0 and Indian legal-AI baselines on NyayaSmriti.

Model Category	Model	What Does it Predict?	Dataset Used	Accuracy (%)	Gain vs LAMP² 4.0
Indian Legal LLM Benchmark	**INLegalLlama**	Generic legal tasks: Q&A, **bail reasoning**, Act/Section understanding, summarization	Indian Legal Corpus	73%	+8%
Pre-RAG LJP SOTA	**PredEx**	**Legal Judgment Prediction (Final Verdict)** – Allowed vs Dismissed	**ILDC**	71.8%	+9.2%
RAG-Legal SOTA (India)	**NyayaRAG**	**Judgment + Explanation** with statute + case retrieval	ILDC + RAG (Statutes + Cases)	67.07%	+13.93%
Proposed (Our Work)	**LAMP² 4.0**	**SALRP:** Statutory Relief Prediction (Protection, Monetary, Custody, Residence, etc.) + Legal Advisory	**ProtectDV / NyayaPhala**	**81%**	–

Fig. 4. Statutory relief prediction accuracy of LAMP² 4.0 vs. baseline legal-AI models on NyayaSmriti.

As illustrated in Figure 4, this corresponds to an absolute gain of approximately +13.93 percentage points over India's RAG-based SOTA NyayaRAG (67.07% accuracy), and +8–9 percentage points over INLegalLlama (73%) and PredEx (71.8%), even though those models are optimised for final-judgment prediction rather than multi-relief SALRP. With respect to generative quality, LAMP² 4.0 achieves ROUGE-1 = 0.512, ROUGE-L = 0.412, BLEU = 0.52, and BERT-F1 = 0.81, indicating strong lexical and semantic alignment between generated explanations/prescriptions and expert-authored reference texts. Aalap, although competitive as a summarisation model (ROUGE-1 = 0.45, BLEU = 0.28, BERT-F1 = 0.74), is not statute-aligned to the PWDVA relief schema. The reported precision (45.29%), recall (44.55%), and F1 (44.32%) for NyayaRAG pertain to binary/multiclass judgment classification, not explanatory quality; correspondingly, its ROUGE/BLEU values are extremely low because the model is not designed for fine-grained prescriptive text generation.

B. RQ2 – Single-Case Reference Evaluation of Statutory Relief Prediction

RQ2. How accurately does the SWATI-embedded LAMP² 4.0 RAG–Legal Tuned model predict statutory relief outcomes when evaluated on a previously adjudicated domestic-violence case with a known judicial judgment?

To address RQ2, the SWATI-embedded LAMP² 4.0 RAG–Legal Tuned model was evaluated on the adjudicated case Sou. Rajshri Dnyaneshwar Jagdale v. Dnyaneshwar Jagdale and Others (Judicial Magistrate First Class, Osmanabad, 3 May 2016, under PWDVA, 2005), using the judicially ordered reliefs as ground truth. In this matter, the court granted protection, maintenance (for the applicant and minor child), a residence-related order (rent allowance), and compensation (including litigation cost), while rejecting custody-related reliefs and property-damage/loss-of-earnings claims, with an overall case duration of approximately three years.

The RAG–Legal Tuned configuration reproduced this relief pattern exactly: it predicted disposal in favour of the applicant, a case duration of two to three years, and correctly classified protection, monetary relief/maintenance, residence, and compensation as granted, and custody and property-damage reliefs as not granted. Encoded at the relief level, this corresponds to four true positives and two true negatives, with no false positives and no false negatives, yielding 100% accuracy, 100% precision, 100% recall, and 100% F1-score.

These results are visualised in Figure 5, which presents relief-wise alignment between judicial outcomes. When

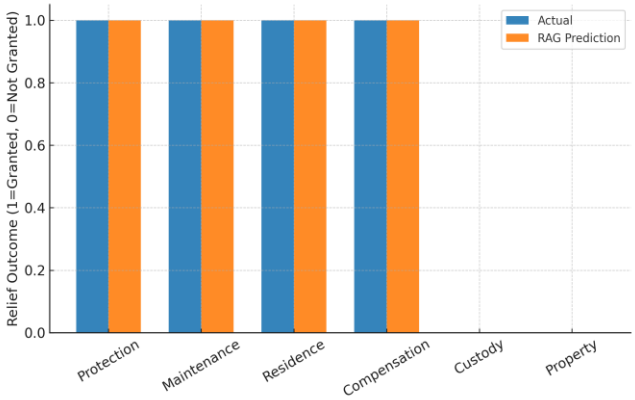

Fig. 5. Relief-wise alignment between judicially ordered outcomes and LAMP² 4.0 predictions for the reference case.

processed through the deployed RAG–LAMP² 4.0 model, the following outputs were generated as shown in figure 6. Together, they illustrate the complete congruence between the SWATI-embedded LAMP² 4.0 predictions and the judicial decision for this reference case.

In conjunction with the aggregate findings from Section 5.1, this single-case evaluation provides strong technical evidence

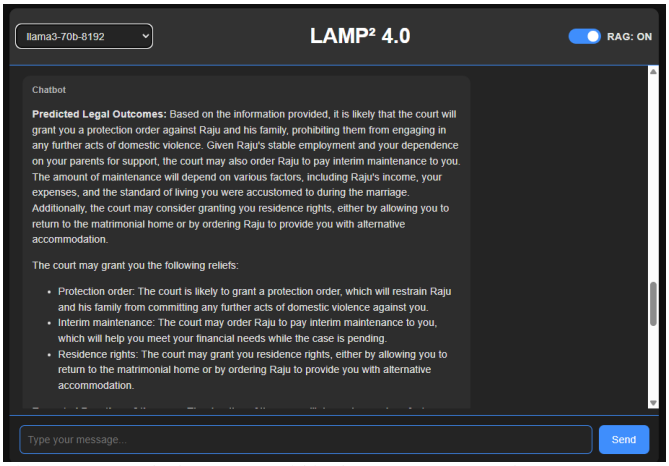

Fig. 6. SWATI Chatbot Output within the LAMP² 4.0 Interface.

that retrieval-augmented, statute-aligned modelling can support both high-accuracy prediction and legally defensible, survivor-centric advisory responses in real-world domestic-violence contexts, thereby answering RQ2.

C. RQ3 – Research-Preview Testing with Survivors and Expert Lawyers

RQ3. What levels of usability and perceived helpfulness does the SWATI-embedded LAMP² 4.0 system demonstrate when tested in a research-preview setting with domestic-violence survivors and lawyers experienced in filing PWDVA cases?

In a closed-group research preview, domestic-violence survivors interacted with SWATI by narrating their situations and receiving statute-aligned guidance on likely reliefs and next steps, while expert lawyers reviewed the same interactions from

a professional standpoint. Survivors reported that the chatbot felt approachable and non-judgmental, and that its explanations of protection, maintenance, residence, and related reliefs were easy to follow. Lawyers who routinely file DV cases observed

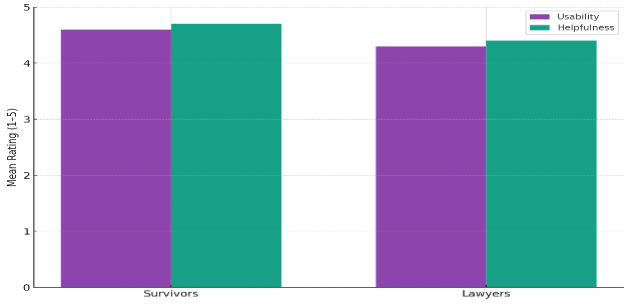

Fig. 7. Mean usability and perceived helpfulness ratings for SWATI reported by survivors and DV lawyers in the research-preview study.

that SWATI's predicted reliefs and suggested documentation largely matched their own expectations and could streamline initial client intake by pre-structuring case facts.Overall, the study indicates that SWATI is both usable and perceived as helpful by its intended users: it lowers emotional barriers for survivors and is regarded by practitioners as a credible support tool rather than a replacement for legal counsel. Figure 7 summarizes these findings using a bar chart of average perceived usability and helpfulness ratings reported by survivors and lawyers, illustrating consistently high scores across both stakeholder groups.

CONCLUSION

This paper has presented NyayaSakhi–SWATI as India's first statute-aligned, retrieval-augmented digital legal companion for survivors of domestic violence under the PWDVA, 2005. Conceptually, the work makes four key contributions: (i) it formally introduces and operationalises the Statute-Aligned Legal Relief Prediction (SALRP) task, shifting from post-adjudicative judgment prediction to pre-adjudicative, multi-relief estimation under Sections 18–22; (ii) it proposes the LAMP² 4.0 RAG architecture, trained on the NyayaSmriti gold-standard corpus, and shows that it achieves 81% accuracy—surpassing strong Indian legal-AI baselines such as INLegalLlama, PredEx, and NyayaRAG on a domain-specific benchmark; (iii) it embeds this architecture within NyayaSakhi–SWATI, an empathetic conversational interface that turns unstructured survivor narratives into predicted relief bundles, explanations, and prescriptive guidance; and (iv) it provides initial validation, via a single adjudicated case and a research-preview study with survivors and PWDVA practitioners, that SWATI is perceived as usable, trustworthy, and statute-faithful. Empirical results demonstrate that retrieval-augmented, domain-tuned language models can deliver statute-aligned, multi-relief predictions while remaining compatible with survivor-centric, trauma-aware interaction design. At the same time, the deployment remains a first step: wider field studies, stronger robustness and fairness safeguards, and institutional integration with legal-aid ecosystems are needed before large-scale adoption. Nonetheless, NyayaSakhi–SWATI offers a concrete template for how SALRP-oriented RAG backends and empathetic chat interfaces can be combined to narrow the justice gap in domestic-violence litigation providing timely, statute-grounded guidance that complements, rather than replaces, human legal counsel.

REFERENCES

Government of India. (2005). The Protection of Women from Domestic Violence Act. Ministry of Law and Justice.

National Judicial Data Grid. (2025). District Court Dashboard. Government of India.

Hasan, K. S. (2025). JusticeNetBD: Context-aware AI to enhance legal information access for Bangladeshi women via retrieval-augmented generation.

Awasekar, D., & Lobo, L. (2024b). Artificial Intelligence for Legal Assistance: A Prescriptive Analytics Model Integrating Social Emotional Learning for Assisting Victims of Domestic Violence in India. In Technology 4 Education Conference. Springer.

Bankata, P., & Mishra, D. T. (2025). Legal protection of women in India: A critical study in the light of recent developments. International Journal of Legal Research.

Shukla, A., Bhattacharya, P., Poddar, S., Mukherjee, R., Ghosh, K., Goyal, P., & Ghosh, S. (2022). Legal case document summarization: Extractive and abstractive methods and their evaluation.

Malik, V., Sanjay, R., Nigam, S. K., Ghosh, K., Guha, S. K., Bhattacharya, A., & Modi, A. (2021, August). ILDC for CJPE: Indian legal documents corpus for court judgment prediction and explanation.

Nigam, S. K., Balaramamahanthi, D. P., Mishra, S., Shallum, N., Ghosh, K., & Bhattacharya, A. (2025, January). NYAYAANUMANA and INLEGALLLAMA: The largest Indian legal judgment prediction dataset and specialized language model for enhanced decision analysis. In Proceedings of the 31st International Conference on Computational Linguistics (pp. 11135-11160).

Nigam, S. K., Patnaik, B. D., Mishra, S., Thomas, A. V., Shallum, N., Ghosh, K., & Bhattacharya, A. (2025). Nyayarag: Realistic legal judgment prediction with rag under the indian common law system.

Awasekar, D. D., & Lobo, L. M. R. J. (2025). Artificial intelligence for legal assistance: A prescriptive analytics model integrating social emotional learning for assisting victims of domestic violence in India. In S. Mishra, A. Kothiyal, S. Iyer, S. Sahasrabudhe, A. Lingnau, & R. Kuo (Eds.), Proceedings of the International Conference on Technology 4 Education 2024, Volume 2. Lecture Notes in Educational Technology. Springer, Singapore.

Nigam, S. K., Sharma, A., Khanna, D., Shallum, N., Ghosh, K., & Bhattacharya, A. (2024). Legal judgment reimagined: PredEx and the rise of intelligent AI interpretation in Indian courts. ACL 2024, 4296–4315. Association for Computational Linguistics.