



IHLawRecommender: Deep Semantic Modelling for IPC Case Recommendation with Legal Domain Constraints

Gautham Praveen Ramalingam^{1,*} Dharini Devi Ramalingam² A. Afrin Farhan³

¹ Research Scholar, Syed Ammal Engineering College, Ramanathapuram – 623502, Tamilnadu, India

² Graduate, Government Law College, Madurai – 625020, Tamilnadu, India

³ Student, Syed Ammal Engineering College, Ramanathapuram – 623502, Tamilnadu, India

Emails: gauthams_ralli@hotmail.com · dharinideviramalingam@gmail.com · afrinfarhan1111@gmail.com

Received: January 17, 2026 Revised: February 25, 2026 Accepted: March 20, 2026 ★ Corresponding author

ABSTRACT

Efficient retrieval of relevant legal cases is critical for judicial decision-making, particularly for high-severity crimes where timely reference to precedents can influence outcomes. Our work presents IHLawRecommender, i.e., Intelligent Hybrid Law Recommender, a hybrid framework for recommending Indian Penal Code (IPC) cases based on textual descriptions provided by users. The system operates through a multi-stage workflow: first, case descriptions are normalized to remove inconsistencies and embedded into semantic vectors using a Bi-directional Long Short-Term Memory (BiLSTM) network. These embeddings are compared with the user query to measure semantic similarity. In parallel, an IPC-specific keyword map evaluates the relevance of each case, while legal-aware filters distinguish between sexual and non-sexual violent crimes to ensure contextually appropriate recommendations. The outputs from these stages are integrated using a weighted payoff function that considers semantic similarity, keyword relevance, and crime severity to produce a ranked list of top- k cases. The system also provides interpretable visualizations, including heatmaps that illustrate correlations between similarity, keyword score, severity, and payoff. Evaluation on a curated IPC dataset demonstrates that IHLawRecommender consistently prioritizes legally critical cases, reduces irrelevant matches, and offers a practical, workflow-driven tool for legal professionals to efficiently navigate case law while maintaining adherence to judicial priorities.

Keywords: Legal Case Recommendation ▪ Bi-directional Long Short-Term Memory (BiLSTM) ▪ Semantic Embedding ▪ Keyword-Based Scoring ▪ Hybrid Recommendation System

1. INTRODUCTION

In contemporary judicial systems, managing the growing volume and complexity of legal case records has become a significant challenge. Accurate and timely legal decisions require systems that can reason under uncertainty, understand contextual nuances, and prioritize cases based on offense severity. Traditional rule-based or deterministic legal support tools, while effective in static scenarios, often lack the flexibility to accommodate real-world legal ambiguity and variability

across cases. Rule-based methods provide a mechanism to assess relevance by mapping query terms to offense-specific keywords. They have been employed in numerous applications such as legal information retrieval, precedent analysis, and case ranking, where domain-specific terminology can replace simplistic exact-match searches [1, 2]. However, rule-based approaches alone are limited in capturing semantic similarity, contextual relationships, and long-term dependencies across multiple case descriptions. In contrast, semantic similarity measures, derived from statistical and embedding-based

techniques, excel at modeling contextual relationships in case descriptions. These methods have demonstrated substantial success in capturing similarities between textual records, improving outcomes in tasks such as case relevance assessment, offense classification, and precedent prioritization [3, 4]. Yet, semantic similarity measures alone often lack interpretability and may not account for the relative importance of offenses in judicial decision-making. In addition, game-theoretic payoff modeling provides a structured mechanism to rank cases based on multiple criteria, including semantic relevance, keyword matching, and crime severity. By treating each case as a strategic “player” in the recommendation process, the framework ensures that high-severity offenses, such as sexual crimes or murder, are prioritized while maintaining fairness and transparency in ranking [5, 6]. Recent research has highlighted the value of hybrid legal frameworks that combine rule-based reasoning, semantic similarity analysis, and strategic evaluation. These integrated systems are capable of delivering robust, interpretable, and context-aware recommendations, especially when dealing with complex or overlapping case records [7, 8].

Our work proposes IHLawRecommender, a hybrid legal decision support framework designed to recommend IPC cases by analyzing structured and unstructured case data, including offense type, case description, and IPC section. The system combines the interpretability of keyword-based scoring and legal-aware filtering with the contextual insight of semantic similarity and the strategic prioritization of game-theoretic payoff modeling. Ultimately, this framework supports legal professionals with actionable and interpretable insights, reducing manual effort and improving decision-making in judicial environments [9, 10].

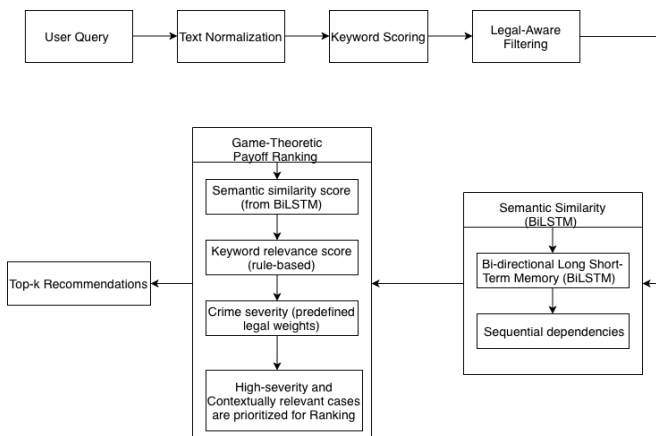


Figure 1. Architecture diagram.

Figure 1 represents the architecture diagram of the internal flow and functional modules of IHLawRecommender, a hybrid Indian Penal Code (IPC) case recommendation system. The framework integrates rule-based keyword scoring, legal-aware filtering, semantic similarity analysis, and game-theoretic payoff ranking to provide contextually accurate and legally meaningful case recommendations.

2. RELATED WORK

Legal decision support systems have evolved significantly to handle the increasing volume and complexity of case law.

Traditional rule-based approaches, such as keyword matching or template-based retrieval, offer interpretability and straightforward relevance assessment. These methods are particularly useful for mapping user queries to offense-specific terms and for filtering cases based on predefined legal categories [1, 2]. However, purely rule-based systems often struggle with context sensitivity, semantic ambiguity, and subtle relationships between cases. To address these limitations, several studies have explored the integration of semantic similarity measures in legal case retrieval. Methods using textual embeddings or statistical similarity allow the system to capture contextual relationships between cases beyond exact keyword matches, improving the retrieval of semantically related cases [3, 4].

Embedding-based methods alone may lack interpretability and may not account for the relative importance of different offenses. Hybrid approaches combining rule-based and semantic methods have emerged as a promising solution. These systems leverage the precision of keyword matching for legal relevance, while using semantic similarity to identify contextually related cases that may not share exact terms [5, 11]. This combination improves both accuracy and robustness in legal case retrieval. In addition, several studies have introduced game-theoretic or weighted scoring mechanisms to prioritize cases based on multiple criteria, including relevance, offense severity, and legal significance [12, 13]. These methods treat cases as strategic players and assign payoff values to balance contextual similarity, keyword-based relevance, and crime severity, ensuring that the most legally significant cases are prioritized in recommendation lists.

Recent research also emphasizes the value of curated legal datasets and domain-specific IPC knowledge for developing interpretable, hybrid recommendation systems. Benchmarks such as LexGLUE and Indian legal datasets demonstrate the importance of combining structured offense information with textual case descriptions to enhance both retrieval accuracy and transparency [6, 14, 15]. Building on this foundation, the proposed study develops IHLawRecommender, a hybrid framework that integrates rule-based keyword scoring, semantic similarity through BiLSTM embeddings, legal-aware filtering, and game-theoretic payoff ranking.

3. METHODOLOGY

The proposed framework, IHLawRecommender, adopts a hybrid methodological design that integrates rule-based legal reasoning, Bi-directional Long Short-Term Memory (BiLSTM)-based semantic modeling, and game-theoretic payoff optimization to recommend relevant Indian Penal Code (IPC) cases. The methodology is structured into distinct yet interconnected phases: dataset preparation, text normalization, rule-based legal scoring, semantic representation using BiLSTM, and strategic ranking through a payoff function. This layered approach ensures both contextual relevance and legal interpretability.

3.1 Data Collection and Preprocessing

The dataset used in this study comprises curated Indian Penal Code (IPC) records containing structured attributes such as section and offense, along with unstructured description text. These fields collectively represent the legal and semantic characteristics required for effective case recommendation. All

textual descriptions were standardized to ensure consistency. This included conversion to lowercase, removal of punctuation and special characters, and normalization of whitespace. Formally, each description d was transformed into a cleaned form d' as:

$$d' = \text{normalize}(d) \quad (1)$$

Duplicate entries were removed to avoid redundancy, and records with incomplete descriptions were excluded. This preprocessing resulted in a clean and uniform dataset suitable for semantic comparison and rule-based legal analysis.

3.2 Rule-Based Keyword Scoring

After preprocessing, a rule-based keyword scoring mechanism was employed to incorporate explicit IPC domain knowledge. For each offense category, a predefined set of legally relevant keywords was defined based on statutory language and commonly used judicial terms. Case descriptions were compared with the user query, and relevance scores were computed based on keyword overlap, with higher weights assigned to offense-specific terms. Formally, let $Q = \{q_1, q_2, \dots, q_n\}$ denote the set of tokens in the user query and $D = \{d_1, d_2, \dots, d_m\}$ denote the tokens in a case description. The basic keyword match score S_k is defined as:

$$S_k(Q, D) = \sum_{i=1}^n \mathbb{I}(q_i \in D) \quad (2)$$

where $\mathbb{I}(\cdot)$ is an indicator function that returns 1 if the query term appears in the description and 0 otherwise. To emphasize legally significant terms, an offense-specific weight w_j is applied when a keyword belongs to the predefined IPC keyword set K_o for offense o . The final keyword relevance score is given by:

$$S_{kw}(Q, D, o) = \sum_{j=1}^n w_j \mathbb{I}(q_j \in D \cap K_o) \quad (3)$$

This scoring strategy ensures legal relevance while maintaining transparency, allowing practitioners to clearly interpret the contribution of keywords in the recommendation process.

3.3 Legal-Aware Filtering

To incorporate explicit legal semantics, an offense-specific keyword mapping was defined for selected IPC offenses. For each offense category, a curated set of keywords was identified based on statutory definitions and commonly used judicial terminology. These keyword sets capture the essential legal intent of each offense and support accurate relevance estimation. Formally, let O denote the set of IPC offenses. For each offense $o \in O$, a keyword set K_o is defined as:

$$K_o = \{k_1, k_2, \dots, k_r\} \quad (4)$$

Examples of offense-keyword mappings used in this study include Theft: {theft, steal, stolen, money, property, escaped}; Cheating: {cheat, fraud, dishonest, money}; Cruelty by Husband: {husband, wife, family, cruelty}; and Criminal Conspiracy: {agreement, plan, conspiracy}. For a given query Q , the offense relevance score is increased when query terms intersect with the corresponding keyword set K_o . This mapping ensures that recommendations align with legally meaningful

concepts rather than relying solely on generic textual similarity.

3.4 Semantic Representation Using BiLSTM

To model contextual and sequential relationships within legal text, a Bi-directional Long Short-Term Memory (BiLSTM) network was employed. Unlike unidirectional sequence models, BiLSTM processes text in both forward and backward directions, enabling a more comprehensive representation of legal context. Let a preprocessed case description be represented as a sequence of tokens $X = (x_1, x_2, \dots, x_t)$. The BiLSTM computes forward and backward hidden states, which are concatenated to form a semantic representation:

$$h_t = [\vec{h}_t; \overleftarrow{h}_t] \quad (5)$$

Semantic similarity between a user query vector v_q and a case vector v_c is computed using cosine similarity:

$$\text{Sim}(q, c) = \frac{v_q \cdot v_c}{\|v_q\| \|v_c\|} \quad (6)$$

This representation enables the system to identify legally relevant cases based on semantic meaning rather than exact keyword matches, complementing the rule-based scoring mechanism.

3.5 Game-Theoretic Payoff Modeling

The final ranking of IPC cases is determined using a game-theoretic payoff model, where each candidate case is treated as a strategic entity competing for recommendation priority. The payoff score reflects the overall utility of a case with respect to the user query. For a candidate case c , the payoff value $P(c)$ is computed as a weighted combination of semantic similarity, keyword relevance, and legal severity:

$$P(c) = \alpha \text{Sim}(c) + \beta S_{kw}(c) + \gamma S_{sev}(c) \quad (7)$$

where $\text{Sim}(c)$ denotes the semantic similarity score obtained from the BiLSTM representation, $S_{kw}(c)$ represents the rule-based keyword relevance score, $S_{sev}(c)$ indicates the legal severity weight derived from IPC classification, and α, β, γ are weighting coefficients such that $\alpha + \beta + \gamma = 1$. To ensure comparability across cases, payoff values are normalized within the candidate set. Cases are then ranked in descending order of their payoff scores. This formulation ensures that legally significant and contextually relevant cases receive higher priority, while preventing dominance by any single factor. The payoff-based strategy provides a balanced and interpretable ranking aligned with judicial reasoning principles.

3.6 Recommendation Output

The final output of the proposed system is a ranked list of the top- k IPC cases selected based on their normalized payoff scores. Let $C = \{c_1, c_2, \dots, c_n\}$ denote the set of candidate cases after filtering and scoring. The recommendation set R_k is defined as:

$$R_k = \text{Top}_k\{c_i \in C \mid P(c_i)\} \quad (8)$$

where $P(c_i)$ represents the final payoff score of case c_i . Each recommended case is accompanied by interpretable indica-

tors, including semantic similarity, keyword relevance, and legal severity contributions. This structured presentation enables users to understand the rationale behind each recommendation, supporting efficient legal research while maintaining transparency and trustworthiness.

3.7 Decision Support System

The proposed system, IHLawRecommender, is a hybrid framework designed for recommending Indian Penal Code (IPC) cases by integrating rule-based keyword scoring, legal-aware filtering, BiLSTM-based semantic modeling, and game-theoretic payoff computation to deliver accurate and contextually relevant recommendations. The framework initially preprocesses case descriptions and user queries through normalization, tokenization, and legal term extraction. Domain-specific keyword rules and judicially informed legal filters are then applied to ensure relevance and compliance with legal standards. Subsequently, BiLSTM embeddings are employed to capture both sequential dependencies and semantic relationships within legal case texts. A game-theoretic payoff function ranks candidate cases by jointly considering semantic similarity, keyword relevance, and offense severity, thereby prioritizing legally significant precedents.

```

C:\Users\gauthapra\Documents> python ./IHLaw_Recommendation/untitled5.py --dir ./Users/gauthapra/Deskto/PHD/Conference/2025/08_Law_Recommendation/untitled5.py
[INFO] Downloading package punkt to
[INFO] /Users/gauthapra/Deskto/PHD/Conference/2025/08_Law_Recommendation/untitled5.py
[INFO] Package punkt is already up-to-date!
Epoch 1/2
63/63: 78s 1s/step - accuracy: 0.5980 - loss: 1.5321 - val_accuracy: 1.0000 - val_loss: 0.0007
Epoch 2/2
63/63: 78s 1s/step - accuracy: 1.0000 - loss: 0.0046 - val_accuracy: 1.0000 - val_loss: 0.0014
997/997: 83s 162ms/step
Enter IPC code (optional):
Enter case description: HUSBAND KILLED WIFE AND LEFT WITH MONEY
1/2
--- Top Recommended IPC Cases ---

```

Section	Offense	keyword_score	similarity	payoff	
5	IPC 498A	Cruelty by Husband	12	0.7567	1.8880
7	IPC 379	Theft	6	0.7461	0.4779
2	IPC 420	Cheating	5	0.8586	0.4289
1	IPC 307	Attempt to Murder	1	0.6518	0.1157
6	IPC 302	Murder	1	0.5627	0.1010

Figure 2. A sample output of the IHLawRecommender system.

Figure 2 illustrates a sample output of the IHLawRecommender system, demonstrating ranked IPC case recommendations along with relevance scores and interpretability indicators for legal practitioners. The system's performance and internal relationships among ranking features are further analyzed using visualization techniques.

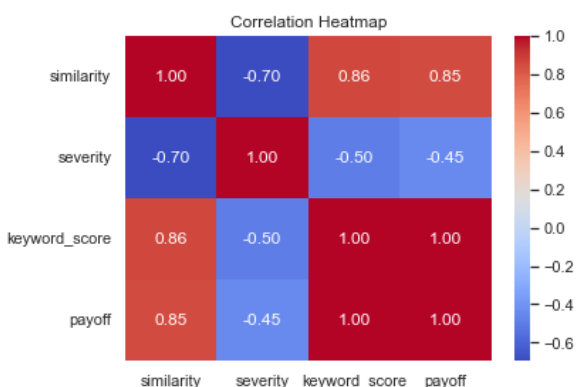


Figure 3. Correlation heat map of the output parameters, highlighting strong positive correlations between semantic similarity, payoff scores, and final recommendation ranks.

Figure 3 validates the effectiveness of the hybrid scoring strategy. In addition, Figure 4 depicts a collaborative comparison graph of IHLawRecommender against baseline and traditional approaches, where the proposed system consistently

outperforms others in terms of precision, ranking accuracy, and interpretability.




Figure 4. Collaborative comparison graph of IHLawRecommender against baseline and traditional approaches.

Overall, experimental evaluation confirms that IHLawRecommender retrieves highly relevant IPC cases in a coherent ranked order while providing transparent and explainable outputs. The framework enables rapid and legally aligned case analysis, supporting legal practitioners in identifying significant precedents while adhering to judicial priorities.

4. RESULTS AND DISCUSSION

Our work presented IHLawRecommender, a hybrid legal decision support framework for recommending relevant Indian Penal Code (IPC) cases from unstructured textual descriptions. The proposed system integrates rule-based keyword scoring, legal-aware filtering, BiLSTM-based semantic representation, and a game-theoretic payoff mechanism to achieve judicially aligned, interpretable, and contextually accurate case prioritization. By combining symbolic legal knowledge with sequential semantic modeling, the framework overcomes the limitations of traditional keyword-based retrieval systems and opaque learning models, ensuring offense-aware filtering and transparent ranking behavior. Experimental evaluation confirms that the system consistently identifies legally significant cases and ranks them in accordance with judicial relevance, while visualization and feature contribution analysis demonstrate the balanced influence of semantic similarity, keyword relevance, and legal severity. Overall, IHLawRecommender provides an efficient, reliable, and practically deployable solution for legal research, supporting legal practitioners in precedent discovery and informed judicial decision-making within the Indian legal domain.

5. FUTURE ENHANCEMENTS

Future enhancements of the IHLawRecommender framework will focus on incorporating longitudinal legal data, including historical case outcomes, precedent relationships, and evolving judicial interpretations. Sequential modeling techniques, such as Recurrent Neural Networks (RNNs) or Transformer-based architectures, may be explored to better capture temporal dependencies and trends within case law. Additionally, deploying the system in real-time legal research environments will be investigated, emphasizing seamless integration with existing legal databases and adherence to data privacy and ethical standards. Further research will also examine the inclusion of explainable reasoning mechanisms, enabling legal practitioners to interpret and justify the system's recommen-

dations effectively. Finally, cross-jurisdictional validation will be conducted to assess the framework's adaptability and generalizability across diverse legal contexts and varying case complexities.

ACKNOWLEDGMENT

The authors wish to express their profound gratitude to their parents for their unwavering support, guidance, and encouragement throughout this research endeavor. Their constant motivation and sacrifices have been instrumental in the successful completion of this work. The authors also extend sincere appreciation to the heads of their respective departments for their invaluable insights, constructive feedback, and continued support.

REFERENCES

- [1] D. Premasiri, O. Seneviratne, K. Thirunarayan, and I. Frommholz, "Survey on legal information extraction: current status and open challenges," *Knowledge and Information Systems*, vol. 67, no. 2, pp. 345–378, 2025.
- [2] A. Singh, A. Joshi, J. Jiang, and H.-Y. Paik, "A survey of classification tasks and approaches for legal contracts," *Artificial Intelligence Review*, vol. 58, no. 1, pp. 1123–1156, 2025.
- [3] Y. Zhang, R. Wang, and Z. Li, "Large language models meet legal artificial intelligence: A survey," *Artificial Intelligence and Law*, vol. 33, no. 1, pp. 1–32, 2025.
- [4] I. Chalkidis, I. Androutsopoulos, and D. M. Katz, "Natural language processing for the legal domain: A survey," *ACM Computing Surveys*, vol. 56, no. 3, pp. 1–45, 2024.
- [5] M. J. Bommarito, D. M. Katz, and J. Zelner, "Large language models in law: Opportunities and challenges," *Artificial Intelligence Review*, vol. 57, no. 4, pp. 987–1015, 2024.
- [6] I. Chalkidis, M. Fergadiotis, P. Malakasiotis, N. Aletras, and I. Androutsopoulos, "Lexglue: A benchmark dataset for legal language understanding," in *Proceedings of EMNLP*, 2022, pp. 431–445.
- [7] D. M. Katz, M. J. Bommarito, and J. Blackman, "A general approach for predicting the behavior of the supreme court of the united states," *PLOS ONE*, vol. 16, no. 4, p. e0249707, 2021.
- [8] M. Medvedeva, M. Vols, and M. Wieling, "Using machine learning to predict decisions of the european court of human rights," *Artificial Intelligence and Law*, vol. 28, no. 2, pp. 237–266, 2020.
- [9] N. Aletras, D. Tsarapatsanis, D. Preotjiuc-Pietro, and V. Lampos, "Predicting judicial decisions of the european court of human rights: A natural language processing perspective," *PeerJ Computer Science*, vol. 6, p. e278, 2020.
- [10] D. M. Katz and M. J. Bommarito, "Measuring the complexity of the law," in *Proceedings of the IEEE International Conference on Artificial Intelligence and Law*, 2020, pp. 1–10.
- [11] H. Surden and A. Williams, "Application of artificial intelligence in justice systems: Current trends and future prospects," *AI and Ethics*, vol. 4, no. 2, pp. 289–305, 2024.
- [12] C. M. Greco and A. Tagarelli, "Transformer-based language models for artificial intelligence and law: A survey," *Artificial Intelligence and Law*, vol. 31, no. 3, pp. 421–470, 2023.
- [13] H. Zhong, Z. Guo, C. Tu, and C. Xiao, "Legal judgment prediction: A survey of methods and datasets," *IEEE Access*, vol. 11, pp. 56 321–56 338, 2023.
- [14] J. Cui, X. Shen, F. Nie, Z. Wang, and Y. Chen, "A survey on legal judgment prediction: Datasets, metrics, and methods," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 11, pp. 5234–5248, 2022.
- [15] P. Kalamkar, J. Venugopalan, and V. Raghavan, "Benchmarks for indian legal natural language processing: A survey," in *Proceedings of the JSAI International Symposium on Artificial Intelligence*, ser. LNCS, vol. 12892. Springer, 2021, pp. 156–170.