

ILSIC: Corpora for Identifying Indian Legal Statutes from Queries by Laypeople

Shounak Paul* Raghav Dogra* Pawan Goyal Saptarshi Ghosh

Indian Institute of Technology, Kharagpur

{shounakpaul95, raghav25t}@kgpian.iitkgp.ac.in,

{pawang, saptarshi}@cse.iitkgp.ac.in

Abstract

Legal Statute Identification (LSI) for a given situation is one of the most fundamental tasks in Legal NLP. This task has traditionally been modeled using facts from court judgments as input queries, due to their abundance. However, in practical settings, the input queries are likely to be informal and asked by laypersons, or non-professionals. While a few laypeople LSI datasets exist, there has been little research to explore the differences between court and laypeople data for LSI. In this work, we create **ILSIC**, a corpus of laypeople queries covering 500+ statutes from Indian law. Additionally, the corpus also contains court case judgements to enable researchers to effectively compare between court and laypeople data for LSI. We conducted extensive experiments on our corpus, including benchmarking over the laypeople dataset using zero and few-shot inference, retrieval-augmented generation and supervised fine-tuning. We observe that models trained purely on court judgements are ineffective during test on laypeople queries, while transfer learning from court to laypeople data can be beneficial in certain scenarios. We also conducted fine-grained analyses of our results in terms of categories of queries and frequency of statutes.

1 Introduction

In the legal domain, laws (statutes) are considered to be one of the most fundamental sources of knowledge that guide principles of jurisdiction (Joshi et al., 2023). In practice, a legal practitioner when faced with a legal situation, typically uses their experience and knowledge to identify the applicable statutes, which can be a time-consuming activity. Hence, the task of automated Legal Statute Identification (LSI) has gained prominence.

Early LSI methods relied on statistical or shallow ML models over small datasets (Salton and

Buckley, 1988; Zeng et al., 2007), but deep learning spurred the creation of large-scale datasets from court judgments, where case facts serve as queries and cited statutes as gold-standard outputs (Xiao et al., 2018; Chalkidis et al., 2021; Paul et al., 2022; Chalkidis et al., 2019; Zhong et al., 2018; Wang et al., 2019). Other sources, like bar exam questions, have also been used (Li et al., 2023b, 2024). These, however, differ linguistically from real laypeople queries, which are informal, ambiguous, and lack legal terminology (Su et al., 2024).

Recent work addresses this gap by building datasets from queries by citizens or laypeople - STARD for Chinese (Su et al., 2024), BSARD/BiBSARD for Belgian law (Louis and Spanakis, 2021; Lotfi et al., 2024), and ALQAC for Vietnamese law (Nguyen et al., 2023; Do et al., 2024). Yet, no study has directly compared judgment-based vs. laypeople-query datasets for LSI; Su et al. (2024) only reported limited indirect gains by using laypeople data with LLMs in other downstream tasks.

To address this research gap, we make the following contributions in this resource paper:

- We develop **ILSIC** (Indian Legal Statute Identification Corpus), a novel collection of datasets based on both laypeople queries as well as court judgment facts for Indian law. Our corpus consists of two independent datasets –
 - (i) **ILSIC-Lay**: A dataset containing 8K laypeople queries covering 567 Indian statutes, meant to serve as the first benchmark for laypeople LSI in the Indian domain;
 - (ii) **ILSIC-Multi**: A dataset containing two different groups of train/validation sets (one based on court facts, the other based on laypeople queries) and a single test set of laypeople queries, all covering a common set of 399 Indian statutes. This corpus allows researchers to directly compare between models being trained on laypeople vs. court data, or any combination of these.

* Equal Contributions

- We run benchmarking experiments on ILSIC-Lay with traditional retrieval methods, such as BM25 (Robertson et al., 2009), SentenceBERT (SBERT) (Reimers and Gurevych, 2019) and SAILER (Li et al., 2023a), as well as state-of-the-art Large Language Models (LLMs) such as GPT-4.1, Llama-3 and Gemma-3, using various experimental setups/variations such as zero and few-shot prompting, retrieval-augmented generation (RAG) and supervised fine-tuning (SFT). Our results demonstrate the relative difficulty of the LSI task, with either GPT-4.1 or fine-tuned versions of the other LLMs not crossing 35% μ F1.
- We utilize the ILSIC-Multi corpus to compare between court facts vs. laypeople queries as input. We devise multiple settings, such as Zero-Shot Inference and Supervised Fine-tuning on either court or laypeople data. We observe that pure court-tuning is completely inefficient for answering laypeople queries. We also created some synthetic laypeople queries by processing court facts with an LLM; this setting can be thought to be somewhere between pure court facts and real-life laypeople queries. However, there is no improvement over pure court data in this setting. We also conducted a transfer learning experiment by sequentially training on court data first followed by laypeople data, which showed limited improvement for only Llama-3.
- We conduct extensive analyses of the above results, by analyzing the model performances across different frequencies of statutes and query categories. We observe that all models have reduced performance over the rare statutes, while the correlation between F1-score and query category frequency is not as strong.

All the datasets, codes and trained model parameters are available at <https://github.com/Law-AI/ilsic>.

2 Related Work

LSI using formal legal documents: Traditionally, court case judgments have been used for creating LSI datasets for countries that follow some degree of the civil law system. Such case judgments contain the facts, i.e., the necessary background information which serves as the input, as well as references to the statutes considered relevant to the case by legal professionals

which serves as the gold-standard output. Xiao et al. (2018) introduced the CAIL2018, the first large-scale dataset for Chinese criminal law, comprising of over 2.6M samples and 183 statutes. Chalkidis et al. (2019) introduced the ECHR dataset based on case facts from the European Court of Human Rights, which involved a binary prediction of whether any statute has been violated. A newer version of the dataset involved multi-label classification over 10 ECHR statutes (Chalkidis et al., 2021). Paul et al. (2022) introduced the ILSI dataset for identifying 100 criminal statutes from the Indian Penal Code. While these datasets were created using court judgment documents, the COLIEE Task 3 family of datasets (Li et al., 2023b, 2024) require retrieving the relevant statutes from Japanese law given questions from the Bar exam.

LSI using informal queries: Although using formal legal texts such as judgment documents can be used as a proxy for modeling LSI, in the practical setting, the automated method should be able to process queries/facts written in an informal way, by laypeople or non-professionals. Lately, some works have focused on creating LSI datasets from such laypeople queries. Su et al. (2024) released STARD, a Chinese LSI dataset based on real-world legal consultation data. Louis and Spanakis (2021) introduced BSARD, a French LSI dataset based on Belgian law, consisting of legal queries by Belgian citizens. This work was later extended to a bilingual setting, through the addition of queries in Dutch (Lotfi et al., 2024). The ALQAC Task 1 family of datasets (Nguyen et al., 2023; Do et al., 2024) require retrieving the relevant statutes from Vietnamese law, given description of legal scenarios in Vietnamese and Thai language. For the Indian legal scenario, Bhattacharya et al. (2019) introduced a very limited set of 50 real-world legal scenarios, annotated with statutes from the Indian Penal Code. Among all these studies, only Su et al. (2024) have made a comparison between formal and informal queries, demonstrating that utilizing laypeople data as the knowledge source of RAG can improve performance of LLMs on other downstream tasks. However, no study, to the best of our knowledge, has conducted a direct comparison between using laypeople queries vs. court queries for the task of LSI.

Dataset	Jurisdiction	Input Type	#Queries	#Statutes
Prior Datasets using formal queries				
ILSI (Paul et al., 2022)	India	Case-facts from judgments	66k	100 IPC statutes
CAIL2018 (Xiao et al., 2018)	China	Criminal case judgments	2.68M	183 Chinese criminal articles
COLIEE Task 3 (Li et al., 2023b, 2024)	Japan	Bar exam-style queries	996 train / 101 test	768 Japanese laws
ECHR (Chalkidis et al., 2021)	EU	Case-facts from ECHR judgments	11k	10 ECHR statutes
Prior Datasets using informal/laypeople queries				
STARD (Su et al., 2024)	China	Laypeople queries from legal consultations	1543	55k Chinese articles
BSARD (Louis and Spanakis, 2021)	Belgium	Legal queries by citizens	1100	22.6k Belgian laws
ALQAC (Nguyen et al., 2023; Do et al., 2024)	Vietnam	Synthetic legal scenarios crafted by laypeople	9846	Various
ILSIC (This work)				
ILSIC-Lay	India	Legal queries by laypeople	8k	569 Indian statutes
ILSIC-Multi	India	Legal queries by laypeople + case-facts from Indian judgments	7k	399 Indian statutes

Table 1: A brief comparison of our ILSIC datasets with prior works

3 Corpora for Legal Statute Identification

As discussed in Section 1, the objective of this study is two-fold: (i) create the first LSR dataset for Indian law that are based on real-life queries asked by laypeople (non-professionals), and (ii) compare the effects of using laypeople queries vis-a-vis queries created from more authentic legal sources such as court judgment documents. To this end, we create two separate, but related corpora:

(i) **ILSIC-Lay**: Consists of a total of 8127 queries ranging across a set of 569 statutes $ILSIC-Lay_{stat}$. All the queries in train, validation and test splits are sourced from real-life queries asked by laypersons on an online legal forum. This is the general dataset for performing LSI with laypeople queries, and consists of a range of statutes, including very rare ones.

(ii) **ILSIC-Multi**: Consists of two distinct groups of train and validation sets, and a single test set based on laypeople queries. The entire corpus spans a set of 399 statutes $ILSIC-Multi_{stat}$, which is a subset of $ILSIC-Lay_{stat}$. One group of train and validation sets, namely $ILSIC-Multi_{train}^{lay}$ and $ILSIC-Multi_{val}^{lay}$ are subsets of $ILSIC-Lay_{train}$ and $ILSIC-Lay_{val}$ respectively, and are to be used for training models with laypeople queries. The other group of train and validation sets, namely $ILSIC-Multi_{train}^{court}$ and $ILSIC-Multi_{val}^{court}$, are based on facts extracted from court judgment documents. The sole test set $ILSIC-Multi_{test}$

consists of laypeople queries and is a subset of $ILSIC-Lay_{test}$. The purpose of this suite of datasets is to evaluate the effectiveness of using laypeople vs. court queries to train models over the same set of statutes, when the end-task is to identify statutes from real-life laypeople queries.

A brief snapshot of our corpus statistics are presented in Table 2.

3.1 Construction of ILSIC-Lay

The Indian law forum <https://kaanoon.com> provides facilities for common users (laypersons) to post any legal query, which is then answered by different lawyers associate with the platform. Each such conversation occurs in the form of a thread, with the user posting their query and lawyers replying to them. The user can ask follow-up queries as well. We scraped $\sim 30k$ such conversations from the website to build the laypeople query datasets.

Statute Extraction and Normalization: Our source of the gold-standard statutes for each asked query comes from the different lawyer responses. Some of the user queries themselves may mention a statute, such as, “I got charged under IPC 279 after crashing my car and injuring some pedestrians ...”, and we leave filter those out using regular expressions. For the rest of the queries, we developed a prompting pipeline with GPT-3.5 Turbo to extract all statute references from the lawyer answers (Prompt provided in App. Table 8). Since lawyers may reference the

Data Split	#Queries	#Words	#Citations
ILSIC-Lay			
ILSIC-Lay _{train}	6465	180.71	2.20
ILSIC-Lay _{val}	826	184.69	2.32
ILSIC-Lay _{test}	836	180.78	2.40
ILSIC-Multi			
ILSIC-Multi _{train} ^{lay}	5793	184.62	2.14
ILSIC-Multi _{val} ^{lay}	735	187.20	2.29
ILSIC-Multi _{train} ^{court}	12930	711.47	2.39
ILSIC-Multi _{val} ^{court}	1652	747.38	2.59
ILSIC-Multi _{test}	757	183.10	2.33

Table 2: Brief Statistics of ILSIC data splits

same statute in different ways, such as “Section 185 of the Motor Vehicle Act” or “s.185 MV Act” (and many other such variations), we had to make extensive use of regular expressions and fuzzy string matching to normalize the statute names to the most formal one (in the above case, it would be “Section 185 of the Motor Vehicles Act, 1988”). By following the above process, we ended with a total set of 600+ statutes.

Query Cleaning and Filtering: Some queries did not have statute references in the respective lawyer answers, and we filtered out all such queries. Some queries were not in English (either containing characters from non-English alphabets, or transliterated into English from other languages), and these were skipped too. After these steps, we had a total of $\sim 12k$ queries spanning 594 statutes, where each query had at least one statute citation in the corresponding lawyer answers. We did a final round of filtration, such that finally each statute had at least three supporting queries. This gave us the final dataset of 8k queries and 567 statutes. For each of the queries, we used the OpenNyAI model (Kalamkar et al., 2022) to detect and mask named entities, to prevent biases for downstream models.

Dataset Quality Validation: We asked a Law expert (a Master in Law graduate from a reputed Law school in India) to manually identify all statutes mentioned in 50 randomly sampled documents. Then we compared the set of statutes identified by the Law expert and the set of statutes identified using the regex approach. Through this manual verification of 50 documents, we confirmed that the regex approach successfully captured over 95% of statute references. The few missed ones were due to non-standard abbreviations, typing errors, and so on. This demonstrates that the coverage of our

extraction pipeline is dependable. While a small fraction of references may not be captured due to formatting inconsistencies, we believe their impact on overall evaluation remains negligible.

Dataset Splits and Statistics: For our train/validation/test splits, we wanted to ensure that some statutes are held-out in the validation and test sets, i.e., some statutes are not cited at all in the train set gold-standard. We randomly selected some statutes as held-out, and tried to find combinations of train/validation/test splits such that the intended ratio between the three splits is approximately 8:1:1 respectively. After multiple iterations, we found a certain combination that satisfied all our criteria. This gave us split sizes of 6465 (train), 826 (validation) and 836 (test), citing a total of 535, 367 and 358 statutes respectively. 8 different statutes are held-out in the validation and test sets each. The queries comprise a total of 10 different legal topics, such as Business Law, Family Law, Taxation, and so on, as tagged by the source law forum (see Table 6 in App. A). Some salient statistics of our dataset are summarized in Table 2.

3.2 Construction of ILSIC-Multi

To conduct analyses regarding the impact of using laypeople vs. court queries for training models, we need to create training corpora for the two different types of queries that are based on the same candidate statute space. We firstly scraped $\sim 50k$ case judgments from the Supreme Court of India and several prominent High Courts. Thereon, we found out the common statutes between those cited across ILSIC-Lay and these court judgments. This gives us ILSIC-Multi_{stat}, a set of 399 statutes that are common to both.

Laypeople Query Components: To create ILSIC-Multi_{train}^{lay}, we basically filtered out those queries from ILSIC-Lay_{train} that cite at least one statute from ILSIC-Multi_{stat}, giving us a total of 5793 queries. We repeat the same process over the validation and test sets of ILSIC-Lay to give us ILSIC-Multi_{val}^{lay} and ILSIC-Multi_{test}, containing 735 and 757 queries respectively.

Court Query Components: For preparing the court-based queries, we wanted to create equivalent train and validation sets as ILSIC-Multi_{train}^{lay} and ILSIC-Multi_{val}^{lay}, but based on court judgment facts. To reflect the real-world scenario where court judgment documents are found in abundance

compared to laypeople queries, we sampled 12930 court documents for ILSIC-Multi^{court}_{train}, which is the double the size of ILSIC-Multi^{lay}_{train}, and spans the same set of 399 statutes. We follow the same procedure for creating ILSIC-Multi^{court}_{val}. For extracting only the fact parts of these judgments, and performing anonymization, we use the OpenNyiAI model (Kalamkar et al., 2022). The court queries are significantly longer as compared to laypeople queries, as seen in Table 2. A detailed analysis of the linguistic differences between court and laypeople data is presented in App. A. A clear sample demonstrating the difference between court and laypeople data is presented in App. C.

4 Methods and General Experimental Settings

As mentioned in Section 1, the goal of this study is to benchmark different models on the ILSIC-Lay dataset, and perform comparisons between court and laypeople queries with the ILSIC-Multi dataset. In this Section, we describe the general models and settings used for the different kind of experiments.

4.1 Retrieval-based Methods

We used multiple retrieval-based methods for benchmarking.

BM25 (Robertson et al., 2009): A sparse information retrieval method based on keyword matching.

SBERT (Dettmers et al., 2023): A dense retrieval approach that uses a BERT-based encoder trained on sentence similarity, to encode queries and statute definitions into a shared semantic embedding space, which are then compared using dot product (<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>).

SAILER (Li et al., 2023a): Another dense retrieval model with a BERT-based encoder that has been pre-trained over legal documents and fine-tuned for legal retrieval (https://huggingface.co/CSHaitao/SAILER_en_finetune).

Experimental Setup: These retrieval methods return, for a given query q , a ranked list of all the candidates (statutes) based on semantic similarity with q . To compare with other approaches, we consider only the top k statutes for each query to be relevant predictions, and the optimal choice of k is determined based on best performance over the respective validation sets.

4.2 Generative Methods

We also benchmarked several modern generative Large Language Models (LLMs) on our dataset.

GPT: We used OpenAI’s GPT-4.1 (<https://platform.openai.com/docs/models/gpt-4.1>), one of the latest and most powerful versions of GPT, for inference experiments. This is one of the latest and most powerful GPT models.

LLaMA (Dubey et al., 2024): LLaMa is one of the most popular open-source LLMs deployed heavily in recent research. We used Llama-3 (8B) (<https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>) under different experimental setups listed below.

Gemma (Kamath et al., 2025): Another popular open-source model, we used Gemma-3 (12B) (<https://huggingface.co/google/gemma-3-12b-it>).

Qwen-3: We used Qwen-3 (4B) (<https://huggingface.co/Qwen/Qwen3-4B>) for inference experiments. Qwen-3 is a recent multilingual LLM with strong instruction-following capabilities.

Phi-4 Mini: We used Phi-4 Mini (4B) (<https://huggingface.co/microsoft/phi-4-mini-instruct>) which is a compact model designed for efficient instruction-following.

Experimental Setups: We use the LLMs in different experimental setups, which we describe below.

Zero-Shot Inference (ZSI): Models were given the prompt to return the most plausible statutes (Section number with Act name) for each query. Most modern LLMs already know most of the statute definitions, which we manually verified for 30% of the most rarely cited statutes.

Two-Shot Inference (TSI): For a given query, we selected two similar training queries based on SentenceBERT (SBERT) similarity from the same category. These examples were then provided in the prompt along with their corresponding gold-standard statutes.

Retrieval-Augmented Generation (RAG): In this setup, we wished to augment the prompt with statute information of some possibly relevant statute candidates. In the setup used in this paper, only statute titles are provided. To identify the possibly relevant statutes, we first used SBERT to rank the training queries similar to fetching two-shot examples (described above), and collected the gold-standard statutes cited by these queries to get

a set of 15 statutes. This approach is an indirect method to obtain possibly relevant statutes, since in Indian law, similar queries are likely to cite similar statutes.

Supervised Fine-tuning (SFT): We also performed supervised fine-tuning using QLoRA (Dettmers et al., 2023) with open-source LLMs.

All generative models, during inference, require a verbalizer that can match the LLM output with the set of target statutes in our dataset. We employ a Verbalization Algorithm for LLM Inference that standardizes free-form outputs into canonical *Section-Act* pairs using hybrid fuzzy matching to ensure consistent and comparable statute references across models. (see App. D for details).

The hyper-parameter settings for the different experiments and the prompts used for the LLMs are listed in App. B and App. E, Table 12 respectively.

4.3 Evaluation Metrics

We evaluated all the experiments under a standard multi-label classification setting, where each query may be associated with multiple applicable statutes. We report both macro- and micro-averaged Precision, Recall, and F1-score. While macro-averaging gives equal weight to each statute, micro-averaging aggregates prediction counts across statutes, and is thus dominated by the frequent statutes.

5 Experiments on ILSIC-Lay

We benchmark all the models described above on ILSIC-Lay and report the results in Table 3.

5.1 Results

The retrieval models, namely BM25, SBERT and SAILER, perform poorly for laypeople queries. Their reliance on formal terms clashes with the informal language in queries, leading BM25 to miss matches and SAILER (trained on US/Canadian text) to fail in generalization. SBERT, though trained on general corpora, performs relatively better. The optimal k values were 5, 3 and 6 for BM25, SBERT and SAILER respectively. Overall, retrieval methods remain weak baselines, compared to generative LLMs which surpass retrieval approaches even in ZSI. GPT-4.1 dominates across metrics, particularly recall (40.91% μ R in ZSI), despite no task-specific tuning. All models improve with TSI, but GPT-4.1 remains ahead in both precision and recall.

Model	Precision		Recall		F1	
	m	μ	m	μ	m	μ
Retrieval-based Methods						
BM25	2.81	2.81	7.79	6.05	3.81	3.84
SBERT	8.05	8.05	13.94	10.24	9.19	9.02
SAILER	4.37	4.37	13.45	10.92	6.03	6.24
Zero-Shot Inference						
GPT-4.1	13.84	25.75	<u>19.51</u>	<u>40.91</u>	14.41	31.60
Llama-3	7.77	15.17	9.40	22.23	7.04	18.03
Gemma-3	8.98	20.77	8.32	22.28	6.94	21.50
Qwen-3	4.88	8.04	5.89	16.73	3.97	10.86
Phi-4 Mini	5.48	16.79	4.94	11.10	4.22	13.37
Two Shot Inference						
GPT-4.1	15.52	33.10	17.26	35.43	14.83	34.23
Llama-3	9.42	17.96	11.56	27.32	9.05	21.67
Gemma-3	10.83	29.87	10.31	29.86	9.58	29.87
Retrieval Augmented Generation						
GPT-4.1	11.83	28.34	16.30	37.92	12.53	32.44
LLaMA3	8.56	23.06	7.83	24.84	7.30	23.92
Gemma3	10.60	25.78	13.54	34.57	10.63	29.54
Supervised Fine-tuning on ILSIC-Lay_{train}						
Llama-3	12.11	31.69	12.41	35.76	10.91	31.52
Gemma-3	11.29	35.34	10.98	31.28	9.95	33.19

Table 3: Performance of all methods on the ILSIC-Lay dataset. All results are in percentage, and all models have been evaluated on ILSIC-Lay_{test}. Best values for **GPT**, **Llama** and **Gemma** are marked with the respective colours, and the best values overall are underlined.

Among open-source models, Llama-3 benefits most from supervised fine-tuning, reaching μ F1 of 31.52%, but performs poorly with RAG, indicating sensitivity to noisy retrieval. Gemma-3 is more balanced; TSI boosts its mF1, RAG lifts μ R (31.53%), and SFT yields its strongest overall trade-off (μ F1 33.19%). Thus, Gemma approaches GPT-4.1 on frequent statutes when fine-tuned, while Llama relies heavily on SFT to stay competitive.

Interestingly, RAG underperforms TSI across all models, suggesting that a small number of in-context examples is more effective than retrieved statutes. This is largely due to limitations of the SBERT retriever used in RAG (for detailed analysis, see App. H). Comparing TSI and SFT, the former favors rare statutes (higher mF1), while the latter biases toward frequent ones (higher μ F1). GPT-4.1 is unique in maintaining smoother performance across both regimes, underscoring its generalization ability even without fine-tuning.

In summary, *GPT-4.1 performs the best off-the-shelf, while Gemma-3 offers the most reliable open-source alternative when fine-tuned, and Llama-3*

improves substantially only with SFT.

5.2 Analysis

To make a more fine-grained analysis of the results in Table 3, we analyze the effect of query category and statute frequency on the μ F1 scores of the best performing models, i.e., GPT-4.1 (ZSI and TSI), and the SFT versions of Llama-3 and Gemma-3.

Effect of Query Category: We divide all the queries in ILSIC-Lay_{test} into their respective categories, and calculate the μ F1 score over each individual category. These results are plotted in Figure 1, where the X-axis represents each category, and the categories are sorted left to right in descending order of frequency over ILSIC-Lay_{test}.

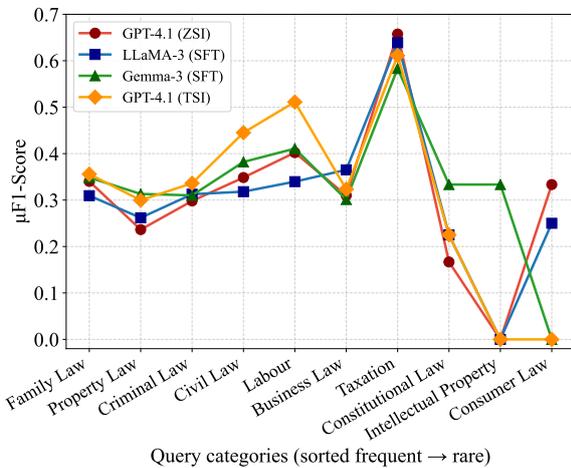


Figure 1: F1 Scores vs query categories sorted according to frequency over ILSIC-Lay_{test} (left to right in descending order of category frequency)

Scores generally decline for rarer categories such as Constitutional, I.P., or Consumer Law. An exception is *Taxation*, where all models achieve unusually high μ F1 (~ 0.6 vs. ~ 0.4). This arises from data distribution: of six *Taxation* statutes in the test set, five also appear in training, leaving only one held-out statute, unlike adjacent categories (e.g., *Business* and *Constitutional Law* with 8 and 3 held-out statutes). Thus, the spike reflects overlap in statute coverage rather than true model capability.

Across categories, Gemma generally outperforms Llama (notably in *Constitutional* and *I.P. Law*), except in *Business*, *Taxation*, and *Consumer Law*, while GPT-4.1 tracks closely with either model depending on the category, showing particular gains in frequent categories under TSI.

Effect of Statute Frequency: We sort all the statutes cited in the gold-standard of

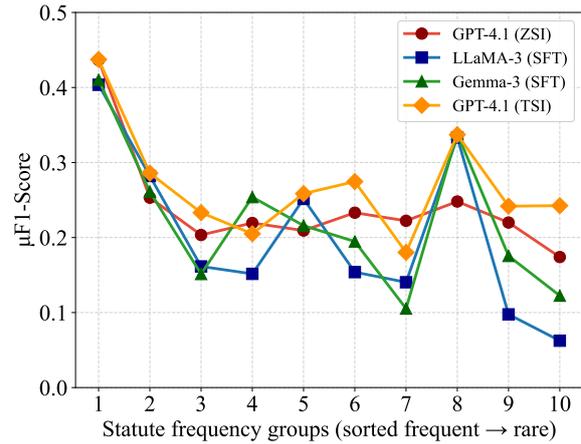


Figure 2: F1 Scores vs Statutes sorted according to frequency over ILSIC-Lay_{test} and divided into 10 groups (left to right in descending order of statute frequency)

ILSIC-Lay_{test}, arrange them into 10 groups and calculate the μ F1 score over each individual group of statutes. These results are plotted in Figure 2, where the X-axis represents each group, and the statute groups are sorted left to right in descending order of frequency over ILSIC-Lay_{test}. Performance decreases as statutes become rarer, except for Group-8, which shows a sudden spike. These statutes mostly align with frequent categories (including *Taxation*), enabling models to exploit indirect category–frequency correlations. Fine-tuned models amplify this effect (steep Group-8 jump), while GPT-4.1 exhibits smoother performance trends, matching fine-tuned models on frequent Groups 1–5 but clearly surpassing them on rare Groups 6–10 (except Group-8), with TSI giving it a further edge on the rarest statutes.

Performance on held-out statutes: As described in Section 3.1, some statutes in the test set are held-out, i.e., only cited by test set queries, and not in train/dev sets. We test the generalization capabilities of the different models by calculating the metrics over the held-out statutes, and observe that GPT-4.1 still outperforms all other methods, while Gemma performs better than Llama in most settings (for more details, see App. G).

Ablation study on RAG pipeline: We conduct a set of ablation experiments on the RAG pipeline by varying the retrieval source, contextual information, and evidence count, and observe that RAG performance is sensitive to these design choices (details in App. I).

Qualitative error analysis: We analyze common

Model	ILSiC-GT		Annotated-GT		ILSiC-A-GT	
	mF1	μ F1	mF1	μ F1	mF1	μ F1
Zero Shot Inference						
Llama-3	1.74	19.33	1.56	11.49	0.79	9.58
Gemma-3	1.82	22.02	1.76	13.46	1.02	13.07
GPT-4.1	3.88	28.97	4.86	22.71	2.80	18.46
Supervised Fine-tuning						
Llama-3	2.70	29.66	2.23	14.99	0.98	13.94
Gemma-3	2.47	32.88	2.17	14.58	1.13	17.00

Table 4: Performance (in %) of base and fine-tuned models on human-annotated ground truth sets.

failure modes across ZSI, TSI, and SFT using representative examples (see App. J).

5.3 Evaluation with Human Experts

To strengthen the reliability of the gold-standard annotations, we conducted a human evaluation with the help of an LL.M. graduate from a reputed Law school in India, hired as a subject matter expert for this project. The expert annotated relevant statutes for 50 randomly sampled queries. We compare three ground truth variants: **ILSiC-GT** (original dataset annotations), **Annotated-GT** (human expert annotations), and **ILSiC-A-GT** (their union). On average, ILSiC-GT contained 2.40 relevant statutes per query, Annotated-GT 5.46, and ILSiC-A-GT 6.53. This indicates a clear progression in coverage, where the original dataset captured only the most essential statutes, while human annotation included additional supporting ones.

When evaluated against the human annotations, models showed slightly lower overall F1 scores due to a reduction in recall, as the expert annotated more statutes than the original dataset. However, *the relative performance trends across the models were consistent* across all ground truth sets, suggesting that the dataset and the expert annotations are well in agreement. Even models fine-tuned on ILSiC continued to outperform base versions when tested on Annotated-GT and ILSiC-A-GT, underscoring that the improvements are not due to overfitting but reflect genuine generalization. Together, these findings demonstrate the robustness of our evaluation framework and dataset reliability.

6 Experiments on ILSiC-Multi

The ILSiC-Multi dataset enables a fair comparison of court vs. laypeople queries over the same set of statutes. We experiment with the following:

Model	Precision		Recall		F1	
	m	μ	m	μ	m	μ
Zero Shot Inference						
GPT-4.1	16.71	28.53	21.93	41.92	17.01	33.95
Llama-3	7.92	25.19	6.04	17.62	5.64	20.74
Gemma-3	10.45	28.10	7.58	23.49	7.35	25.59
Supervised Fine-tuning on ILSiC-Multi^{lay}_{train}						
Llama-3	12.67	29.11	14.74	36.50	12.21	32.39
Gemma-3	12.74	36.87	13.29	32.88	11.96	34.76
Supervised Fine-tuning on ILSiC-Multi^{court}_{train}						
Llama-3	7.27	17.62	10.52	22.98	6.52	19.95
Gemma-3	9.74	10.42	10.42	23.38	8.05	23.10
Supervised Fine-tuning on ILSiC-Multi^{syn}_{train}						
Llama-3	8.05	16.74	10.55	22.52	7.11	19.20
Gemma-3	8.71	19.78	10.25	23.21	7.50	21.36
Transfer learn (ILSiC-Multi^{court}_{train} \rightarrow ILSiC-Multi^{lay}_{train})						
Llama-3	13.62	29.87	15.52	36.56	13.06	32.88
Gemma-3	12.03	36.05	13.00	33.40	10.99	34.67

Table 5: Performance of all methods on the ILSiC-Multi dataset. All results are in percentage, evaluated on ILSiC-Multi_{test}.

- (i) **Zero-Shot Inference (ZSI):** Baseline inference.
- (ii) **Laypeople Query Supervised Fine-tuning:** Models trained on ILSiC-Multi^{lay}_{train} and validated on ILSiC-Multi^{lay}_{val}, reflecting the test set.
- (iii) **Court Query Supervised Fine-tuning:** Models trained/validated on ILSiC-Multi^{court}_{train} and ILSiC-Multi^{court}_{val}, to test efficacy of court facts when evaluated on laypeople queries.
- (iv) **Synthetic Query Supervised Fine-tuning:** Court queries in ILSiC-Multi^{court}_{train} and ILSiC-Multi^{court}_{val} are re-written into laypeople style using GPT-4o-mini (Prompt in Table 13), giving ILSiC-Multi^{syn}_{train} and ILSiC-Multi^{syn}_{val}.
- (v) **Transfer Learning:** Sequential fine-tuning on ILSiC-Multi^{court}_{train}/ILSiC-Multi^{court}_{val} followed by ILSiC-Multi^{lay}_{train}/ILSiC-Multi^{lay}_{val}.

6.1 Results

Table 5 shows results of different experiments on ILSiC-Multi. The trends mirror the results over ILSiC-Lay (Table 3), with GPT-4.1 yielding the best mF1 under ZSI, but fine-tuned Llama-3 and Gemma-3 surpass it in μ F1. Court-only training is ineffective, giving worse μ F1 than ZSI and only marginal mF1 gains, confirming the linguistic mismatch between court and laypeople queries. Synthetic laypeople queries behave similarly, failing to capture laypeople-specific language. Transfer

learning improves Llama-3 (best Llama result overall) but gives Gemma-3 only modest gains.

6.2 Analysis

We repeat the analyses of Section 5.2 for query categories and statute frequency. Pure laypeople training consistently outperforms pure court training, while transfer learning results fluctuate (see App. F).

7 Conclusion and Future Work

In this resource paper, we create a new corpus **ILSIC** for identifying Indian legal statutes from laypeople queries. This is the first corpus that allows researchers to experiment with direct comparison between court facts vs. laypeople queries for LSI over the same set of statutes (ILSIC-Multi dataset). Our benchmarking over the ILSIC-Lay dataset shows that fine-tuned Llama and Gemma models are able to somewhat match the ZSI performance of GPT-4.1. In the comparison experiments over ILSIC-Multi, we observe that training purely over a large court dataset is completely ineffective for inference over laypeople queries, as compared to training on a smaller laypeople dataset. Transfer learning from court to laypeople data can help in certain cases, but there are still a lot of gaps to be addressed in terms of the relationship between court and laypeople data for LSI. In future, we would like to try out other ways of combining the court and laypeople training data to further improve upon training purely on a small laypeople dataset.

Acknowledgments

We thank the anonymous reviewers and the metareviewer for their insightful comments and suggestions. We also thank the law expert *Shubham Tiwari*, Master of Laws (LLM) from the Rajiv Gandhi School Of Intellectual Property Law, India, who helped us conduct the annotation study.

This research is partially supported by the IIT Mandi iHub and HCl Foundation (iHub) (through the project titled “Large Language Model for Legal Assistance”) and by the IIT Kharagpur Technology Innovation Hub on AI for Interdisciplinary Cyber-Physical Systems (AI4ICPS) (through the project titled “NyayKosh: Multilingual Resources for AI-based Legal Analytics”).

Limitations

In this paper, we created a new dataset for LSI over laypeople queries in the Indian legal scenario. We are the first, to the best of our knowledge, to create a dataset that allows direct comparison between using court facts vs. laypeople queries for LSI. However, there are a few limitations of our study, which we describe briefly in this section.

Since most legal forums in India do not provide public access to question-answer data, we were able to collect laypeople queries from only a single website. Also, the dataset is limited to the English language, and does not cover the high linguistic diversity of India. It remains an important future work to develop a multilingual dataset of laypeople queries in various Indian languages.

For benchmarking on the ILSIC-Lay dataset, we have considered a limited number of models. This limited model coverage is partly due to hardware constraints. All experiments are conducted on a single 48GB Nvidia RTX A6000 GPU, which restricts evaluation to models of moderate size and prevents full fine-tuning of larger models. In future, there is a need to experiment with more legal retrieval-focused models.

For comparing between court facts and laypeople queries for LSI, we have experimented with a limited number of settings. In terms of combining court and laypeople data for training, we have only utilized transfer learning, which produces marginal benefits over pure laypeople training only in limited cases. This aspect of the work demands further exploration by using other techniques to combine the two types of training data.

Ethical Considerations

In this work, we develop a dataset that allows for the statute identification from laypeople queries. This task is crucial for the legal domain, and systems for this task are likely to help legal professionals as well as common persons. The AI methods / models are designed only to provide relevant recommendations to the legal professionals, or preliminary guidance to laypeople users, and are *not* expected to be integrated directly into the decision-making process of the judicial system.

The dataset is constructed from *publicly available* webpages on <https://kaanoon.com>, an online Indian law forum. We ensured to collect only those queries which have been kept public by the people who asked the queries. We did not collect

any query that was made private by the person who asked the query. Also, we manually reviewed 50 documents and found almost no personal identifiers such as person-names in the queries. Location names are present, and have been retained for legal relevance (since applicable statutes often depend on the particular Indian state/region), ensuring privacy while preserving essential context.

The dataset is made available specifically for academic, non-commercial research, under the the Creative Commons Attribution-NonCommercial-ShareAlike (CC BY-NC-SA 4.0) license.

References

- Paheli Bhattacharya, Kripabandhu Ghosh, Saptarshi Ghosh, Arindam Pal, Parth Mehta, Arnab Bhattacharya, and Prasenjit Majumder. 2019. Overview of the fire 2019 aila track: Artificial intelligence for legal assistance. In *FIRE (working notes)*, pages 1–12.
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. Neural legal judgment prediction in english. *arXiv preprint arXiv:1906.02059*.
- Ilias Chalkidis, Manos Fergadiotis, Dimitrios Tsaratsanis, Nikolaos Aletras, Ion Androutsopoulos, and Prodromos Malakasiotis. 2021. *Paragraph-level rationale extraction through regularization: A case study on European court of human rights cases*. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 226–241, Online. Association for Computational Linguistics.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. *Qlora: Efficient finetuning of quantized llms*. *arXiv preprint arXiv:2305.14314*.
- Dinh-Truong Do, Son T Luu, Trang Pham, Trung Vo, Nguyen-Hoang Chu, Quang-Huy Chu, Cuong Nguyen, Minh Nguyen, An Trieu, Dat Nguyen, et al. 2024. A summary of the alqac 2024 competition. In *2024 16th International Conference on Knowledge and System Engineering (KSE)*, pages 422–427. IEEE.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- Abhinav Joshi, Akshat Sharma, Sai Kiran Tanikella, and Ashutosh Modi. 2023. *U-CREAT: Unsupervised case retrieval using events extrAcTion*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13899–13915, Toronto, Canada. Association for Computational Linguistics.
- Prathamesh Kalamkar, Astha Agarwal, Aman Tiwari, Smita Gupta, Saurabh Karn, and Vivek Raghavan. 2022. Named entity recognition in indian court judgments. *arXiv preprint arXiv:2211.03442*.
- Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, et al. 2025. Gemma 3 technical report. *CoRR*.
- Haitao Li, Qingyao Ai, Jia Chen, Qian Dong, Yueyue Wu, Yiqun Liu, Chong Chen, and Qi Tian. 2023a. *Sailer: Structure-aware pre-trained language model for legal case retrieval*. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*, page 1035–1044, New York, NY, USA. Association for Computing Machinery.
- Haitao Li, You Chen, Zhekai Ge, Qingyao Ai, Yiqun Liu, Quan Zhou, and Shuai Huo. 2024. Towards an in-depth comprehension of case relevance for better legal retrieval. In *JSAI International Symposium on Artificial Intelligence*, pages 212–227. Springer.
- Haitao Li, Weihang Su, Changyue Wang, Yueyue Wu, Qingyao Ai, and Yiqun Liu. 2023b. *Thuir@coliee 2023: Incorporating structural knowledge into pre-trained language models for legal case retrieval*.
- Ehsan Lotfi, Nikolay Banar, Nerses Yuzbashyan, and Walter Daelemans. 2024. Bilingual bsard: Extending statutory article retrieval to dutch. *arXiv preprint arXiv:2412.07462*.
- Antoine Louis and Gerasimos Spanakis. 2021. A statutory article retrieval dataset in french. *arXiv preprint arXiv:2108.11792*.
- Chau Nguyen, Son T Luu, Thanh Tran, An Trieu, Anh Dang, Dat Nguyen, Hiep Nguyen, Tin Pham, Trang Pham, Thien-Trung Vo, et al. 2023. A summary of the alqac 2023 competition. In *2023 15th International Conference on Knowledge and Systems Engineering (KSE)*, pages 1–6. IEEE.
- Shounak Paul, Pawan Goyal, and Saptarshi Ghosh. 2022. Lesicin: A heterogeneous graph-based approach for automatic legal statute identification from indian legal documents. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 11139–11146.
- Nils Reimers and Iryna Gurevych. 2019. *Sentence-bert: Sentence embeddings using siamese bert-networks*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.

- Gerard Salton and Christopher Buckley. 1988. [Term-weighting approaches in automatic text retrieval](#). *Information Processing & Management*, 24(5):513–523.
- Weihang Su, Yiran Hu, Anzhe Xie, Qingyao Ai, Quezi Bing, Ning Zheng, Yun Liu, Weixing Shen, and Yiqun Liu. 2024. [Stard: A chinese statute retrieval dataset derived from real-life queries by non-professionals](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10658–10671.
- Pengfei Wang, Yu Fan, Shuzi Niu, Ze Yang, Yongfeng Zhang, and Jiafeng Guo. 2019. [Hierarchical matching network for crime classification](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR’19*, page 325–334, New York, NY, USA. Association for Computing Machinery.
- Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Yansong Feng, Xianpei Han, Zhen Hu, Heng Wang, et al. 2018. [Cail2018: A large-scale legal dataset for judgment prediction](#). arXiv preprint arXiv:1807.02478.
- Yiming Zeng, Ruili Wang, John Zeleznikow, and Elizabeth Kemp. 2007. [A knowledge representation model for the intelligent retrieval of legal cases](#). *International Journal of Law and Information Technology*, 15(3):299–319.
- Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Chaojun Xiao, Zhiyuan Liu, and Maosong Sun. 2018. [Legal judgment prediction via topological learning](#). In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 3540–3549.

Appendix

Table of Contents

A	Details of the ILSIC Datasets.	13
B	Hyper-parameters for Fine-tuning Experiments	13
C	Court and Laypeople Query Examples on the Same Statute	14
D	Verbalization Algorithm for LLM inference	15
E	Prompts used for LLM Experiments.	15
F	Details of Experiments on ILSIC-Multi	17
G	Performance on Held-Out Statutes	19
H	Analysis of SBERT Retrieval Performance	20
I	Ablation Study on the RAG Pipeline	20
J	Qualitative Error Analysis	21

List of Tables

6	Category-wise query distribution for the ILSIC laypeople dataset	13
7	Readability statistics for the Laypeople and Court in ILSIC-Multi.	13
8	Prompt used for Legal Reference Extraction Task	14
9	Evaluation across held-out statutes under different experimental setups.	15
10	Court Query on Consumer Protection and IPC Section 427	16
11	Laypeople Query on Consumer Protection and IPC Section 427	17
12	Prompt used for Zero-Shot Inference and Supervised Fine-Tuning Task	18
13	Prompt used for Synthetic laypeople dataset generation	18
14	Prompt used for Two-Shot Inference Task	19
15	Prompt used for RAG Task	20
16	SBERT retrieval performance and downstream effects on RAG compared to Zero-Shot Inference (ZSI).	20
17	RAG ablation results with macro (m) and micro (μ) evaluation metrics.	21
18	Example of hallucinated statute predictions in Zero-Shot Inference (ZSI).	21
19	Example of non-applicable statute predictions in Two-Shot Inference (TSI).	22

20	Example of technically misaligned statute predictions in Supervised Fine-Tuning (SFT).	22
----	--	----

List of Figures

3	Comparison of fine-tuning models across different settings and categories of queries	17
4	Comparison of fine-tuning models across different settings and frequency of statutes	18

A Details of the ILSIC Datasets

The ILSIC-Lay dataset comprises queries spanning a broad spectrum of legal categories, with separate splits for training, validation, and testing. It includes both commonly encountered topics such as Family Law and Property Law, as well as specialized areas like Intellectual Law and Taxation. The distribution of queries across categories reflects the diversity and prevalence of legal issues seen in real-world contexts, supporting the development and evaluation of robust legal language models. The detailed category-wise statistics are provided in Table 6

Category	Train	Val	Test
Family Law	2333	258	335
Property Law	1924	251	223
Criminal Law	1065	154	156
Civil Law	645	80	52
Labour	190	42	28
Business Law	172	24	28
Constitutional Law	56	9	4
Taxation	49	5	6
Consumer Law	27	3	2
Intellectual Law	4	0	2
Total	6465	826	836

Table 6: Category-wise query distribution for the ILSIC laypeople dataset

Linguistic Differences between Laypeople and Court data: To understand the linguistic gap between the two domains, we conducted a detailed readability and lexical analysis of the laypeople and Court subsets in the ILSIC-Multi dataset. The results, summarized in Table 7, reveal distinct stylistic and structural contrasts between the two types of texts.

The laypeople subset displays a higher Flesch Reading Ease score (64.53) and a lower Flesch-Kincaid Grade Level (9.29), suggesting that these queries are easier to comprehend and written in a simpler linguistic style. In contrast, the Court subset has a lower Reading Ease (46.81) and a higher Grade Level (12.38), reflecting the formal and complex nature of judicial text.

Lexical diversity: Type-Token Ratio (TTR) captures the proportion of unique words to total words, serving as a measure of vocabulary variety. However, since TTR is sensitive to text length, we also report Herdan’s C, a length-normalized measure of lexical richness. The laypeople subset shows a

Metric	Laypeople	Court
Average Flesch Reading Ease	64.53	46.81
Average Flesch-Kincaid Grade Level	9.29	12.38
Lexical Diversity – TTR	0.0203	0.0081
Lexical Diversity – HerdanC	0.7237	0.7027
Sentence Length Distribution	18.278	20.674

Table 7: Readability statistics for the Laypeople and Court in ILSIC-Multi.

higher raw TTR (0.0203 vs. 0.0081) owing to its shorter, conversational phrasing. Yet, the Herdan’s C values are similar (0.7237 vs. 0.7027), suggesting that while Court texts are more repetitive and formulaic, they maintain comparable normalized lexical diversity.

Sentence length distribution: Average sentence length is slightly higher in Court documents (20.67 words) compared to laypeople queries (18.28 words), reflecting the longer and syntactically denser constructions typical of formal legal writing. The laypeople subset, by contrast, contains shorter and more direct sentences, consistent with natural, query-like expressions.

Overall, these measures demonstrate that the laypeople subset is more readable, lexically varied, and concise, whereas the Court subset is denser, more technical, and syntactically complex. This linguistic divergence underscores the challenge for models to generalize effectively between lay and formal legal domains.

B Hyper-parameters for Fine-tuning Experiments

For all fine-tuning experiments, we tried out for maximum of 5 epochs with early stopping. For SAILER, we used a batch size of 16, and 1 positive sample and 47 negative samples per query. SAILER was fine-tuned using Contrastive Loss with in-batch negative sampling. We used a l.r. $5e-6$. For LLMs, we used QLoRA for fine-tuning, with the LoRA parameters being $\alpha = 32$, $r = 8$, dropout at 0.05, l.r. $1e-5$ and 4-bit quantization. We used a batch size of 4, and trained using the standard Cross Entropy Loss. The maximum sequence length was kept at 32k. The LLMs were trained over *completion only*.

You are a highly knowledgeable legal expert. Your task is to identify and extract all legal references, such as statutes_2993, acts, sections, and legal cases, explicitly mentioned in the input text. Follow these instructions carefully.

Instructions:

- 1) Extract **only** legal statutes_2993, acts, sections, and cases that are **explicitly mentioned** in the input text. Do not infer or assume any references.
- 2) Categorize the extracted information into the following structures:
Sections: Each entry must include both the section number and its explicitly stated act.
Cases: List all legal cases explicitly mentioned, including citations if provided.
- 3) Replace "Indian Penal Code" with IPC and "Criminal Penal Code" with CrPC wherever applicable.
- 4) Ensure high accuracy and completeness. Do not omit or add any information beyond what is explicitly stated.
- 5) Do not output acts separately. Each section must explicitly include its associated act.
- 6) If multiple sections belong to the same act, list them as separate objects.
- 7) If no references are found, return empty lists ([]).
- 8) Handle partial or ambiguous references by including them as-is without interpretation.

Output Format (JSON only):

```
{
  "sections": [
    {
      "act": "IPC or CrPC or other explicitly stated act",
      "section": "section number as mentioned"
    }
  ],
  "cases": [
    "list of explicitly mentioned cases"
  ]
}
```

Rules:

- Do not infer missing acts or sections.
- Maintain official legal naming standards.
- Output must strictly follow the specified JSON structure.
- Do not include commentary, explanations, or extra text.

User Input:

QUERY:
{Target_query_text}

Return only the JSON object.

Table 8: Prompt used for Legal Reference Extraction Task

C Court and Laypeople Query Examples on the Same Statute

To illustrate the differences in language and style between laypeople and court queries relating to the same statute, we present a representative pair of

queries below. Each pair includes a query from a court document and a corresponding laypeople query, both referring to the same legal provision(s).

Example: Section 12 of The Consumer Protection Act, 1986 and Section 427 of The Indian Penal

Model	Precision		Recall		F1 Score	
	m	μ	m	μ	m	μ
Zero Shot Inference (ZSI)						
Llama-3	3.25	77.78	18.75	17.08	22.88	28.00
Gemma-3	5.00	100.00	18.75	17.08	26.05	29.17
GPT-4.1	6.25	100.00	38.55	36.59	44.05	53.58
Two Shot Inference (TSI)						
Llama-3	3.25	80.00	20.83	20.42	30.42	31.38
Gemma-3	50.00	100.00	23.75	21.96	31.75	34.60
GPT-4.1	6.25	100.00	30.21	27.78	40.38	46.29
Retrieval-Augmented Generation (RAG)						
Llama-3	25.00	8.75	6.25	4.17	7.14	4.77
Gemma-3	25.0	100.0	5.63	4.84	9.17	9.41
Supervised Fine-Tuning (SFT)						
Llama-3	12.5	100.0	7.29	7.22	11.31	9.17
Gemma-3	12.5	100.0	5.00	4.88	7.15	9.31

Table 9: Evaluation across held-out statutes under different experimental setups.

Code, 1860

Description:

The court query [Table 10](#) is highly structured, referencing procedural history, evidence, and legal claims in formal and technical language. In contrast, the laypeople query [Table 11](#) tells a personal story in simple language, explains events chronologically, and focuses on seeking practical advice and justice, without legal terminology.

D Verbalization Algorithm for LLM inference

To ensure consistent and interpretable statute references during generation, we employ a structured verbalization pipeline that converts model outputs into standardized *Section–Act* pairs. The process begins by isolating potential statute mentions from the generated text using regular expressions tuned to legal numbering patterns (e.g., “Section 420 of the Indian Penal Code”). Each extracted mention is then normalized through lowercasing, removal of connective fillers (e.g., *of, under, by*), and punctuation stripping.

For robust alignment with canonical references, we perform hybrid matching that combines exact section matching with fuzzy act similarity computed via a weighted ensemble of RapidFuzz metrics, ratio, token-sort, and token-set, supplemented by a Jaccard word-overlap score and selective boosts for key legal terms. A match is accepted only when the section identifier matches exactly and the act-name similarity exceeds a confidence threshold ($\tau \approx 70\text{--}85$).

This verbalization step harmonizes minor surface variations (e.g., “IPC” \leftrightarrow “Indian Penal Code”, “Motor Vehicles Act” \leftrightarrow “MV Act”), yielding consistent statute identifiers across LLM outputs. On a manually validated sample of 50 generation instances, the algorithm achieved approximately 95% correct statute normalization, with most residual discrepancies arising from irregular act abbreviations or nested references within multi-clause sections.

E Prompts used for LLM Experiments

The prompt shown in [Table 15](#) serves as a foundational component in our approach to statutory retrieval and model fine-tuning. By simulating the role of a legal expert, the prompt provides clear and restrictive instructions for the extraction of statutory references, requiring the identification of precise section and act pairs directly relevant to a given legal scenario. This structure not only ensures the completeness and accuracy of extracted information but also standardizes output for consistent model training and evaluation. As a result, the prompt enables effective fine-tuning of language models on legal tasks, supports robust retrieval of statutes from complex legal texts, and lays the groundwork for building practical AI tools for legal research, compliance checking, and legal decision support in the Indian legal context.

The prompt detailed in [Table 13](#) is designed to guide the model in generating clear and accessible summaries of Indian court cases for non-expert audiences. By instructing the model to avoid legal jargon, technical references, or placeholder terms, and to present the facts in a conversational and easily understandable manner, the prompt ensures the creation of realistic laypeople-style case summaries. This approach is essential for constructing synthetic datasets that reflect how actual litigants or the general public might describe their legal issues online. Such datasets are crucial for training and evaluating models intended for laypeople-facing legal applications.

The prompt shown in [Table 14](#) corresponds to the *two-shot inference setup*, where the model is guided using two semantically similar example pairs before generating the output for the target query. Each example provides a query–text and its corresponding list of statute predictions, enabling the model to observe the expected input–output structure. This few-shot formulation helps the model

The present appeals arise out of the impugned judgment and order dated 26.04.2013 in Revision Petition No. 2032 of 2012 and order dated 23.07.2013 in Review Petition No. 253 of 2013 passed by the National Consumer Disputes Redressal Commission, New Delhi (hereinafter referred to as the National Commission), whereby the petitions challenging the order dated 29.02.2012 passed by the Haryana State Consumer Disputes Redressal Commission were dismissed. The brief facts of the case which are required to appreciate the rival legal contentions advanced by the learned counsel appearing on behalf of the parties are stated in brief as hereunder:- The appellant was the owner of a Tata Motors goods carrying vehicle bearing registration No.HR-67-7492. The vehicle was insured with the respondent- Company vide policy No. 15019923334104992 with effect from 31.07.2009, valid upto 30.07.2010. The risk covered in this policy was to the tune of Rs.2,21,153/-. The said vehicle met with an accident on 11.02.2010 on account of rash and negligent driving of the offending vehicle bearing registration no. UP-75-J 9860. In this regard, an FIR No.66 of 2010 dated 11.02.2010 was registered with the jurisdictional Police Station, Sadar, Fatehabad, for the offence punishable under [SECTION], [SECTION], [SECTION] and [SECTION] of the [ACT] (hereinafter referred to as the [ACT]). The appellant incurred expenses amounting to Rs.1,64,033/- for the repair of his vehicle and also informed the respondent- Company about the accident and damage caused to his vehicle. In this connection, the respondent- Company appointed one Mr. [ENTITY], as the Surveyor to assess the damage caused to the said vehicle. After inspecting the vehicle, the Surveyor assessed the damage caused to the vehicle at Rs.90,000/-, whereas the appellant had preferred a claim for a sum of Rs.1,64,033/- with supporting bills. In addition to above, the respondent-Company appointed M/s Innovation Auto Risk Claim Manager for the purpose of investigation. According to the report of the investigator, five passengers were travelling in the goods-carrying vehicle, though the seating capacity of the vehicle as per the registration certificate was only 1+1. On the basis of findings of the said report, the respondent-Company vide letter dated 26.07.2010 rejected the claim of the appellant for the reason that the loss did not fall within the scope and purview of the insurance policy. Aggrieved of the letter of rejection of the claim of the appellant by the respondent-Company, he filed Complaint No.517 of 2010 against the respondent-Company dated 17.09.2010 before the District Consumer Disputes Redressal Forum, Sonapat (hereinafter referred to as the District Forum) under [SECTION] of the [ACT] for the claim of Rs.1,64,033/- towards the repair of his vehicle on the ground that the rejection of the claim amounts to deficiency in service on the part of the respondent-Company. The respondent-Company filed a detailed written statement before the District Forum disputing the claim of the appellant. It took the plea that the complainant had violated the terms and conditions of the policy, as five passengers were travelling in the goods-carrying vehicle at the time of accident, whereas the permitted seating capacity of the motor vehicle of the appellant was only 1+1. The District Forum on the basis of the pleadings of the parties and the materials on record considered the judgment of the National Commission in the case of [PRECEDENT]], wherein it was held that if the number of persons travelling in the vehicle at the time of the accident did not have a bearing on the cause of accident, then the mere factum of the presence of more persons in the vehicle would not disentitle the insured claimant from claiming compensation under the policy towards the repair charges of the vehicle paid by the appellant. The District Forum accordingly directed the respondent-Company to settle the claim of the appellant on non-standard basis upto 75% of the amount spent for effecting repairs to the damaged vehicle after taking into consideration the claim amount of Rs.1,64,033/-. The District Forum further directed the respondent-Company to settle the amount to be paid to the appellant along with interest at the rate of 9% per annum from the date of lodging of the claim by the appellant with the respondent-Company. The respondent-Company was further directed to pay Rs.2,000/- for rendering deficient service, causing mental agony and harassment and towards litigation expenses incurred by the appellant. Aggrieved of the order of the District Forum, the respondent Company preferred an appeal before the State Commission urging various grounds.

Table 10: Court Query on Consumer Protection and IPC Section 427

I am a retired person and I wanted to build a roof top green house garden. As I don't know anyone who is in the field I searched online and found some companies in Indiamart website. I got several names and chose one (but no basis for my selection as the guy called me within 30min). Initially he was very polite and so thought he is nice person ... so I entered the contract with him. Initially I paid him Rs 50000 (cheque) and he brought the structural materials next day and instructed his guys as how to the work. The work went well for a few days and subsequently all went wrong (meaning no work) as he kept giving excuses. While I gave a threatening call he sent another guy (he said his brother) and took up the work to complete the structural work sincerely. Subsequently, the same brother told me that we need to get the roof material and need Rs 50000. I said the guy who made contract with me should tell and then I give, within a few mins the guy called and told me that you give 50000 cheque. As soon as they took money from bank both of their cell is off all time. Some times it rings but no response. I can afford the money though I am retired. But as a true citizen of India I want them to get punished. I would appreciate if you could suggest me to take forward to curb such criminal. Besides, your suggestion of a honest lawyer will help me to make these guys not to cheat anyone in their life. I live in whitefield area in Bangalore> Thanks and I hope that I will get fruitful suggestion.

Table 11: Laypeople Query on Consumer Protection and IPC Section 427

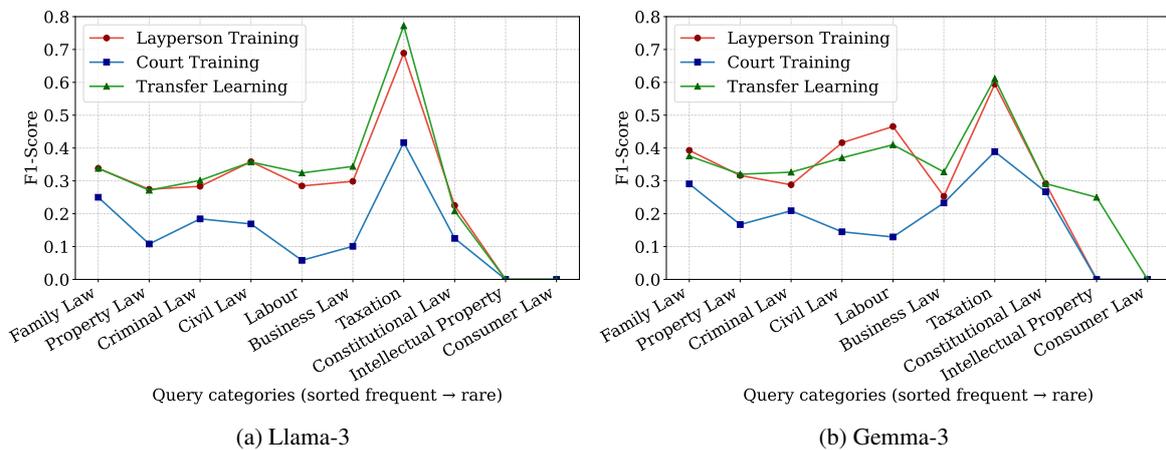


Figure 3: Comparison of fine-tuning models across different settings and categories of queries

generalize the pattern of statute identification by leveraging contextual cues from analogous cases. .

The prompt shown in Table 15 is designed for the retrieval-augmented generation (RAG) stage. Here, the model receives a user query along with a set of candidate statutes retrieved via SBERT similarity and their formal definitions. The model is explicitly instructed to select only the most relevant statutes from this constrained set, adhering strictly to the given candidates without rephrasing or inventing new entries. This structured prompt ensures interpretability with the retrieved legal context, thereby improving reliability in downstream re-ranking and selection tasks.

F Details of Experiments on ILSIC-Multi

Since our objective in this Section is to compare the effects of different fine-tuning data, we plot only the results for pure laypeople training, pure court training and transfer learning for both Llama and Gemma.

Effect of Query Category: The performance distribution with categories for both models are shown in Figure 3. We observe the same spike in performance for *Taxation* for both models as seen in Figure 1. For Llama-3, court training significantly underperforms for all categories, while transfer learning performs marginally better as compared to laypeople training (see Figure 3a). Similar trends are observed for Gemma-3 as well (see Figure 3b), however, transfer learning clearly underperforms compared to pure laypeople training for two categories, *Civil* and *Labour* Law.

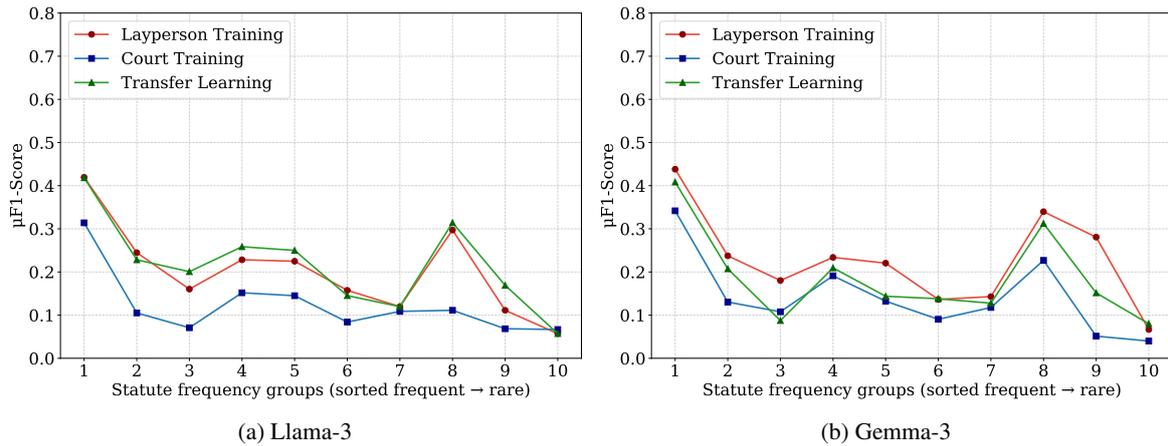


Figure 4: Comparison of fine-tuning models across different settings and frequency of statutes

You are a legal expert on Indian law. Given the text below, identify all applicable legal sections and their full act names. Follow these rules exactly:

1. Identify every relevant provision (section number + full act name) you are 100% certain applies under Indian statutes.
2. Output only one list in square brackets, formatted like: ["Section X of Act Name"; "Section Y of Act Name"; ...]
 - Each entry must be a quoted string with the section first, then the full act name.
 - Each entry needs both section number AND full act name
 - Separate entries with a semicolon.
 - Do not include anything outside this single list.
3. Exclude any entry missing either section or act (no incomplete pairs).
4. Do not repeat identical section-act pairs (no duplicates).
5. Do not add explanations, labels, or extra-only the list itself.
6. Always spell out the full name of each act (no abbreviations).
7. Include only provisions that are clearly and directly relevant (no speculative or uncertain entries).

Table 12: Prompt used for Zero-Shot Inference and Supervised Fine-Tuning Task

You are a legal expert specializing in Indian court cases. I will provide you with the facts of a court case. Your task is to generate a laypeople’s summary of the case, written in a single, continuous paragraph, without using bullet points, numbered lists, or section headers. The summary should be clear, concise, and easy to understand for someone without legal expertise, but it must include all relevant facts and details. Write in a conversational style, as if you are personally describing your situation and seeking advice, similar to how a layperson might post their question on an online forum.

Important: Do not label or mention legal entities, sections, acts, or precedents with placeholders such as [ENTITY], [SECTION], [ACT], [PRECEDENT], etc. Do not add, invent, or fill in any legal names, sections, or references that are not present in the original case facts I provide. If the court facts I provide do not mention a statute (such as a section or act), you must not include any such statute in your output, under any circumstance. Only summarize the information actually present in the input.

Only output the summary in the required style.

Table 13: Prompt used for Synthetic laypeople dataset generation

You are a legal expert on Indian law. Given the text below, identify all applicable legal sections and their full act names. Follow these rules exactly:

1. Identify every relevant provision (section number + full act name) you are 100% certain applies under Indian statutes.
2. Output only one list in square brackets, formatted like: ["Section X of Act Name"; "Section Y of Act Name"; ...]
 - Each entry must be a quoted string with the section first, then the full act name.
 - Each entry needs both section number AND full act name.
 - Separate entries with a semicolon.
 - Do not include anything outside this single list.
3. Exclude any entry missing either section or act (no incomplete pairs).
4. Do not repeat identical section-act pairs (no duplicates).
5. Do not add explanations, labels, or extra text-only the list itself.
6. Below are two EXAMPLES of the task (input 1/2 with output1/2). Then produce the answer ONLY for input 3 by writing output3. Stop immediately after you close the single list with a ‘]’.

input 1: {Example_1_query_text}

query-religion: {Religion_1}

output1: ["Section ... of Act ..."; "Section ... of Act ..."]

input 2: {Example_2_query_text}

query-religion: {Religion_2}

output2: ["Section ... of Act ..."; "Section ... of Act ..."]

input 3: {Target_query_text}

query-religion: {Religion_3}

output3: [Model to generate this list]

Table 14: Prompt used for Two-Shot Inference Task

Effect of Statute Frequency: The performance distribution across frequency of statutes are shown in Figure 4. Again, similar to Figure 2, we observe a spike in performance for both models on Group-8. For Llama-3, transfer learning performance is similar (but slightly less) to pure laypeople training across all groups, while court training underperforms. For Gemma-3, however, transfer learning performance reduces to the level of pure court training for many of the frequent groups (Groups 3-5).

G Performance on Held-Out Statutes

To assess the models’ generalization beyond the training distribution, we evaluated their performance on a subset of 8 held-out statutes and the corresponding 41 queries that referenced them. This analysis captures the models’ ability to reason over unseen statutory provisions rather than relying on memorized associations. As shown in Table 9, the zero-shot inference (ZSI) and two-shot inference

(TSI) settings demonstrated the strongest results among the evaluated setups. In particular, GPT-4.1 achieved the highest overall scores, with macro- and micro-F1 values of 0.4475 and 0.5358 in the zero-shot setting, reflecting consistent generalization to unfamiliar statutes. Gemma and Llama 3 models followed with moderate gains, especially under the two-shot configuration, indicating that limited in-context exemplars improve adaptation to unseen sections.

In contrast, the retrieval-augmented generation (RAG) and supervised fine-tuning (SFT) approaches exhibited lower precision and recall, suggesting that exposure to retrieved or previously seen statutes did not transfer effectively to unseen legal provisions. These results emphasize that the generative reasoning capacity of large language models, particularly GPT-4.1, plays a more decisive role in extrapolating to new statutes than retrieval or task-specific fine-tuning. Overall, the

You are a legal expert on Indian law. Select statutes that best apply to the user’s situation, but **ONLY** from the provided CANDIDATE_STATUTES.

Rules:

- 1) Output exactly one list in square brackets using double quotes, e.g., ["Section X of Act"; "Section Y of Act"].
- 2) Each entry must be a quoted string in the exact same form as provided in CANDIDATE_STATUTES (no rephrasing).
- 3) Do not add anything outside that one list. No commentary.
- 4) Do not invent or add statutes not in CANDIDATE_STATUTES.
- 5) Prefer fewer, more precise provisions over many generic ones.
- 6) If none clearly apply, return [].

User message (per query).

QUERY:

{Target_query_text}

CANDIDATE_STATUTES:

- "Section . . . of Act . . . "
- "Section . . . of Act . . . "
- . . .

STATUTE_DEFINITIONS:

- "Section . . . of Act . . . ": {definition text}
- "Section . . . of Act . . . ": {definition text}
- . . .

Return only the bracketed list.

Table 15: Prompt used for RAG Task

held-out evaluation highlights both the challenge of statute generalization and the potential of few-shot prompting to mitigate it.

H Analysis of SBERT Retrieval Performance

We evaluate the effectiveness of SBERT as the retriever in our RAG pipeline by measuring its ability to retrieve the ground-truth statute within the top-15 candidates. As shown in Table 16, SBERT achieves a macro precision of only 31.99%, indicating limited retrieval quality.

This retrieval bottleneck directly affects the downstream RAG performance, as the language model is conditioned on a noisy and incomplete candidate set. Consequently, RAG does not yield improvements over Zero-Shot Inference (ZSI) and, in some cases, performs slightly worse. For this reason, we compare RAG primarily against ZSI rather than fine-tuned models, since both settings rely on base models without task-specific training.

Overall, these results suggest that the effectiveness of the RAG setup is constrained by the retriever, rather than the reasoning capability of the language model itself.

Setup	Precision (m)	Recall (m)	F1 (m)
SBERT (Top-15)	31.99	21.23	24.36
RAG (LLaMA3)	7.51	13.00	8.46
RAG (Gemma3)	5.70	6.69	5.40
ZSI (LLaMA3)	7.70	9.40	7.04
ZSI (Gemma3)	8.98	8.32	6.94

Table 16: SBERT retrieval performance and downstream effects on RAG compared to Zero-Shot Inference (ZSI).

I Ablation Study on the RAG Pipeline

To better understand the factors affecting the performance of our RAG pipeline, we conduct a series of ablation experiments that systematically modify key components of the retrieval and context construction process. All ablations are evaluated using

Setup	Model	Precision		Recall		F1	
		m	μ	m	μ	m	μ
Base-RAG	LLaMA3	8.56	23.06	7.83	24.84	7.30	23.92
Base-RAG	Gemma3	10.60	25.78	13.54	34.57	10.63	29.54
Stat-RAG	LLaMA3	5.63	9.24	5.34	9.33	4.08	9.29
Stat-RAG	Gemma3	8.94	12.49	11.40	18.05	8.29	14.76
Descr-RAG	LLaMA3	7.51	13.59	13.00	31.76	8.46	19.03
Descr-RAG	Gemma3	5.70	13.82	6.69	17.95	5.40	15.62
Top10-RAG	LLaMA3	9.39	16.97	7.69	20.38	6.21	18.52
Top10-RAG	Gemma3	9.39	22.20	12.53	30.87	9.57	25.83
Top5-RAG	LLaMA3	8.03	22.62	7.57	19.72	6.85	21.07
Top5-RAG	Gemma3	9.27	24.93	9.96	25.30	8.55	25.11

Table 17: RAG ablation results with macro (m) and micro (μ) evaluation metrics.

the same base language models and evaluation protocol as the main RAG setup. The quantitative results for all ablation settings are reported in Table 17.

Ablation on Retrieval Source (Stat-RAG): In the RAG setup reported in the main paper, SBERT is used to retrieve relevant evidence by computing similarity between the test query and the set of training queries, followed by selecting the top-15 gold statutes cited by those training queries. As an ablation, we instead retrieve evidence by directly computing SBERT similarity between the test query and the set of all statutes, and select the top-15 statutes to pass to the model. This setting resembles a classical retrieval and re-ranking setup and helps isolate the effect of the retrieval source.

Ablation on Contextual Information (Descr-RAG): In the default RAG setup, we provide title of each retrieved statute as context. To evaluate the impact of contextual richness, we conduct an ablation where both the title and the full statutory description are provided in the prompt.

Ablation on Evidence Count (Top10-RAG and Top5-RAG): The main RAG experiments use the top-15 retrieved statutes as evidence. To study the effect of evidence size, we additionally evaluate RAG setups that use only the top-10 and top-5 retrieved statutes. This ablation helps assess whether reducing the amount of retrieved evidence mitigates noise introduced by imperfect retrieval.

Overall, these ablations allow us to analyze how retrieval strategy, contextual information, and evidence count individually influence the effectiveness of the RAG pipeline.

J Qualitative Error Analysis

To better understand the failure modes of different modeling approaches, we conduct a qualitative error analysis across Zero-Shot Inference (ZSI), Two-Shot Inference (TSI), and Supervised Fine-Tuning (SFT) using the Gemma model. We categorize errors into three broad types based on their relationship to the gold-standard statutes. For each category, we provide a representative example consisting of the input query, the ground-truth statutes, and the model’s prediction.

Hallucinated Statutes: In some cases, the model predicts statutes that do not exist in the Indian legal corpus. These errors indicate hallucination, where the model generates spurious statute references. We observe this error type primarily in ZSI and TSI settings, while SFT does not produce hallucinated statutes in our experiments. A representative example is provided in Table 18

<p>Query: I have agriculture land at village Ottu, Teh. Rania, Distt. Sirsa, Haryana. I own a 4-acre plot that is not connected to any public path. Access is through an approach path owned by my uncle and another person, who often prevent passage. The path has existed for 35 years. Is there any legal or revenue rule that allows me to use this path or obtain legal access to my land, which is 18 karam away and would require crossing my uncle’s land?</p>
<p>Ground Truth: Section 15, Section 32, Section 35, and Section 4 of the Indian Easements Act, 1882</p>
<p>ZSI Output: Section 125 and Section 13 of the Hindu Marriage Act, 1955</p>

Table 18: Example of hallucinated statute predictions in Zero-Shot Inference (ZSI).

The output is a hallucination: the Hindu Mar-

riage Act, 1955 does not contain “Section 125”, and marriage-related provisions are unrelated to an easement/right-of-way dispute. The correct statutes come from the Indian Easements Act, 1882.

Non-applicable Statutes: In this category, the model predicts statutes that are valid and exist in the legal corpus but are not applicable to the given query and are often unrelated in legal scope. These errors reflect imprecise statute selection, where the model associates a statute with the query despite little or no legal relevance. This error type is observed across ZSI, TSI, and SFT. A representative example is provided in [Table 19](#).

Query: SIR, what can I do if a retired army officer is teasing society members, children, and women? We have filed a police case, but no action has been taken as he has strong connections with officials. The person continues to harass children as well.

Ground Truth: Section 156 of the Code of Criminal Procedure, 1973

TSI Output: Section 354, Section 355, and Section 356 of the Indian Penal Code, 1860

Table 19: Example of non-applicable statute predictions in Two-Shot Inference (TSI).

The predicted IPC sections relate to assault and use of criminal force, which are not applicable to the described acts of verbal harassment and intimidation in this query, and therefore do not align with the legal context of the complaint.

Technically Misaligned Statutes: Here, the model predicts statutes that are legally related to the gold standard but differ in technical applicability or scope, reflecting partial understanding of the query. This error type is observed across ZSI, TSI, and SFT. A representative example is provided in [Table 20](#).

“Section 13 of The Indian Easements Act, 1882” is legally related to the dispute but differs in technical applicability from the gold-standard provisions governing prescriptive easements. “Section 145 of The Code of Criminal Procedure, 1973” concerns preventive proceedings over breach of peace and is not technically applicable to the underlying right-of-way claim.

Overall, this analysis shows that hallucination is mainly associated with prompting-based methods (ZSI and TSI), while SFT errors primarily arise from non-applicable or partially applicable statute

Query: I have agriculture land at village Ottu, Teh. Rania, Distt. Sirsa, Haryana. I own a 4-acre plot that is not connected to any public path. Access is through an approach path owned by my uncle and another person, who often prevent passage. The path has existed for 35 years. Is there any legal or revenue rule that allows me to use this existing path or obtain legal access to my land, which is 18 karam away and would require crossing my uncle’s land?

Ground Truth: Section 15, Section 32, Section 35, and Section 4 of the Indian Easements Act, 1882

SFT Output: Section 13 of the Indian Easements Act, 1882; Section 145 of the Code of Criminal Procedure, 1973

Table 20: Example of technically misaligned statute predictions in Supervised Fine-Tuning (SFT).

predictions. These findings highlight complementary strengths and weaknesses of prompting and supervised approaches for legal statute identification.