Team-SVNIT at JUST-NLP 2025: Domain-Adaptive Fine-Tuning of Multilingual Models for English–Hindi Legal Machine Translation

Rupesh Dhakad and **Naveen Kumar** and **Shrikant Malviya**
Sardar Vallabhbhai National Institute of Technology, Surat
dhakadrupesh666@gmail.com,{naveenkumar,shrikant}@coed.svnit.ac.in

Abstract

Translating sentences between English and Hindi is challenging, especially in the domain of legal documents, due to the specialized legal terminology and the lengthy, complex sentences that often accompany them. In this paper, we fine-tune and compare multiple pretrained multilingual translation models, including the `facebook/nllb-200-distilled-1.3B`, on a corpus of 50,000 English–Hindi legal sentence pairs provided for the shared task. The training pipeline includes preprocessing, context windows of 512 tokens, and decoding methods to enhance translation quality. The proposed method secured **1st place** on the official leaderboard. We obtained the following scores on various metrics: BLEU 51.61, METEOR 75.80, TER 37.09, CHRF++ 73.29, BERTScore 92.61, and COMET 76.36. These results demonstrate that fine-tuning multilingual models for a domain-specific machine translation task enhances performance. Our code is released to the public for further exploration https://github.com/Rupeshdhakad06/JUST-NLP-LMT.

1 Introduction

Legal machine translation is more difficult than general translation. It needs both accurate language modelling and correct handling of legal terms (Panezi and O'Shea, 2023). The JUST-NLP 2025 shared task[1] deals with English–Hindi legal translation. The two languages differ in structure and in the way legal contexts are embedded. Accurate translation is not just about replacing words. It also requires keeping the legal meaning and intent the same across both systems (Way, 2016).

Many problems are explored in this area. Using the same legal terms consistently in all contexts is challenging (Altakhaineh, 2025). The lack of parallel legal data limits the amount of supervised training that can be done (Raja and Vats, 2025). In the legal domain, even a small translation mistake can cause serious problems (Llop, 2025). Recent developments in large-scale multilingual NMT, particularly the No Language Left Behind effort, have yielded strong cross-lingual transfer across nearly 200 languages (Costa-jussà et al., 2022, 2024). However, the applicability of these models to specialized domains, such as legal text—especially for Indian languages—remains relatively underexplored (Nair et al., 2024). Within the Indic NLP community, systems such as IndicTrans (Ramesh et al., 2021) and IndicTrans2 (Gala et al., 2023) have broadened multilingual coverage from 11 to 22 languages. Still, hurdles such as rich morphology, multiple scripts, and code-switching persist and complicate model performance on real-world legal corpora (Suman et al., 2023; Sheshadri and Soman, 2023).

This manuscript describes Team-SVNIT's proposed Legal-MT system and experimental evaluation. Our main contributions are: **(1)** An empirical comparison of five candidate translation systems, which identifies the `facebook/nllb-200-distilled-1.3B` model as the best-performing backbone; **(2)** A preprocessing pipeline designed to clean noisy legal text extracted from the Dataset. **(3)** A training regimen using extended contexts (512 tokens), a cautious learning rate (2e-5), and cosine-based scheduling; **(4)** A top-ranking submission that placed first on the task leaderboard; **(5)** A manual, qualitative appraisal of 100 samples to uncover model strengths and recurring error modes.

2 Related Work

2.1 Neural MT for Indic Languages

IndicTrans (Ramesh et al., 2021) was one of the first large-scale, multilingual NMT efforts for 11 Indian languages, employing language-aware pre-

[1] https://exploration-lab.github.io/JUST-NLP/

Proceedings of the 1st Workshop on NLP for Empowering Justice (JUST-NLP 2025), pages 148–154
December 24, 2025 ©2025 Association for Computational Linguistics

processing and transformer-based architectures. IndicTrans2 (Gala et al., 2023) brought coverage to 22 Indic languages and further refined knowledge distillation methods for scalability. The NLLB (Costa-jussà et al., 2022, 2024) project expanded translation into more than 200 languages with extensive coverage of the Indic families. Distilled versions of these models, 600M and 1.3B parameters, retain impressive translation performance with low computational cost (Koishekenov et al., 2023). The architecture is based on a sparsely gated mixture-of-experts architecture that allows for optimal use of parameters without the computational overhead from dense models.

Despite these advances, some challenges are persistent: the Indic NMT systems have to grapple with morphosyntactic richness, orthographic variations across scripts, limited parallel data, and prevalence of code-switched content (Raja and Vats, 2025; Naveen et al., 2024).

2.2 Legal Domain NMT

Neural machine translation for legal texts differs significantly and is challenging compared to general-domain texts due to the scarcity of domain-specific corpora, specialized terminology, and stringent accuracy demands. Complex syntactic structures in legal texts have to be rendered faithfully, as does the translation of jurisdiction-specific terminology (Way, 2016; Panezi and O'Shea, 2023). Minor mistranslations might have detrimental consequences for legal interpretation and the conduct of proceedings (Llop, 2025).

Altakhaineh et al. show that machine-translated legal content is often fraught with critical semantic and syntactic errors, requiring heavy human post-editing (Altakhaineh, 2025). This emphasizes the importance of domain adaptation, model fine-tuning, and human verification if the systems are to be reliably deployed in a legal context (Princeton, 2025).

2.3 Multilingual Model Fine-tuning

Fine-tuning is a critical step in adapting multilingual models for domain-specific corpora, such as legal text. Cosine annealing with warm restarts helps preserve multilingual prior knowledge and mitigates catastrophic forgetting during training (Loshchilov and Hutter, 2017). Using large batch sizes through gradient accumulation makes training more stable and helps it converge (Han et al., 2024). Parameter-efficient methods, such as LoRA, save resources (Nair et al., 2024). However, full fine-tuning is still better for legal translation, where correctness is more important than speed or cost.

3 Task Description

3.1 Dataset

The JUST-NLP 2025 Legal MT shared task includes an English–Hindi parallel dataset(Singh et al., 2025). It covers different areas of law, such as constitutional, civil, criminal, and administrative. Table 1 shows the statistical details of the dataset. For clarity and to illustrate the nature of the translations, Table 3 in Appendix 4 displays representative English–Hindi sentence pairs selected from the training data.

Split	Pairs	Avg(Eng)	Avg(Hin) in words
Train	50,000	29.3	31.1
Valid	5,000	26.8	30.2
Test	5,000	26.1	–

Table 1: Dataset statistics (average words per sentence).test data translation was not given so mentioned with (–).

The legal sentences in this dataset are long and complex. The average English sentence length in the training data is 29.3 words. The Hindi translations are about 6% longer, averaging 31.1 words. This shows the need for longer context windows during training. The dataset also includes common legal phrases, legal citations, numbers, and some noise from digitization, such as mixed scripts and encoding errors.

4 Dataset Examples

Legal sentences can be challenging to translate. For example, a sentence like *"The appellant, being aggrieved by the judgment and decree dated 15th March 2023 passed by the Hon'ble High Court of Delhi in Civil Appeal No. 2345 of 2022, prefers this present appeal under Section 96 of the Code of Civil Procedure, 1908"* contains numerous legal references, dates, and laws that require careful translation. Latin terms such as *"res judicata"*, *"sub judice"* and *"amicus curiae"* also need proper transliteration and meaning adjustment in Hindi legal language.

Model	Params	Langs	Arch	Features
Helsinki-NLP/opus-mt-en-hi	77M	2	Transformer	Lightweight
facebook/nllb-200-distilled-600M	600M	200	Trans+MoE	Conditional routing
ai4bharat/indictrans2-en-indic-1B	1.0B	22	Transformer	Indic-specialized
law-ai/InLegalTrans-En2Indic-1B	1.0B	Indic	Transformer	Legal domain
facebook/nllb-200-1.3B	1.3B	200	Trans+MoE	Standard version
facebook/nllb-200-distilled-1.3B	**1.3B**	**200**	**Trans+MoE**	**Our choice (distilled)**

Table 2: Model selection prioritizes multilingual capacity, sufficient parameters, and architectural setting.

English	Hindi
according to the learned counsel for the appellant, Allahabad Bank was the owner of the vehicle, as the vehicle in question was unexplained by the Bank.	अपीलार्थी के विद्वान अधिवक्ता के अनुसार इलाहाबाद बैंक वाहन का मालिक था क्योंकि प्रश्नाधीन वाहन बैंक के यहां आडमानित था ।
both these writ petitions are, thus, allowed.	इस प्रकार, ये दोनों रिट याचिकाएं अनुज्ञात की जाती हैं ।
so far as the applicability of Section 20 is concerned, it is a case of trial.	जहां तक धारा 20 के लागू होने का सवाल है, यह विचारण का एक मामला है ।

Table 3: Sample training pairs from the legal English–Hindi parallel corpus.

4.1 Evaluation Metrics

The evaluation utilizes six standard translation metrics to assess the model's performance. These are combined into one score called AutoRank. It is defined as:

$$\text{AutoRank} = \frac{1}{6}\sum_{i=1}^{6} M_{i,\text{norm}} \quad (1)$$

The metrics include BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), TER (inverted) (Snover et al., 2006), CHRF++ (Popović, 2017), BERTScore (Zhang et al., 2019), and COMET (Rei et al., 2020). Each score is scaled between 0 and 100.

5 System Architecture

5.1 Model Selection

We tested five translation models, including standard and distilled ones. Table 2 shows the full comparison.

5.2 Preprocessing

There are usually minor errors in Dataset related to legal texts. Our preprocessing pipeline corrected such errors. It substitutes line breaks with spaces and standardized dashes and quotation marks. It also solves the encoding issues, eliminates English words left behind in Hindi text and minimizes the additional spaces. The measures ensure the text remains neat without distorting legal terms, numbers, and references.

Hyperparameter	Value
Base Model	facebook/nllb-200-distilled-1.3B
Max Seq Len	512 tokens
Epochs	20 (early stop)
Batch Size	32
Gradient Accum	16 steps
Effective Batch	512
Learning Rate	2e-5
Scheduler	Cosine w/ restarts
Warmup Ratio	0.1
Optimizer	AdamW
Precision	FP16

Table 4: Training hyperparameters.

5.3 Training Configuration

Table 4 presents our training setup. A context length of 512 covers 99% of the corpus, ensuring full sentence coverage. A large batch size of 512 enables stable optimization, while a conservative learning rate of 2e-5 helps preserve multilingual representations. Cosine scheduling improves convergence by avoiding local minima, and early stopping prevents overfitting. Using FP16 reduces memory usage by 40%, allowing for larger batches and doubling training speed.

5.4 Inference

During inference, the model employs beam search with a width of 4 to identify the optimal translation. The maximum output length is set to 512 tokens, with an n-gram penalty of 3 to avoid repetition. Early stopping ensures efficient decoding once an end token is reached. The process is deterministic, ensuring consistent results, and runs in batches of 64 for faster translation generation.

Model	Params	BLEU	ROUGE	CHRF
Helsinki-OPUS (Without Training)	77M	24.0	50.2	51.3
NLLB-600M + Fine-tuning	600M	43.2	65.5	61.3
ai4bharat/indictrans2-en-indic-1B + Fine-tuning	1.0B	44.0	68.6	62.8
Helsinki-OPUS + Fine-tuning	77M	46.3	70.1	68.9
law-ai/InLegalTrans + Fine-tuning	1.0B	48.1	68.2	66.5
Facebook/NLLB-1.3B + Fine-tuning	1.3B	50.1	73.1	69.4
Facebook/NLLB-1.3B-distilled + Fine-tuning	**1.3B**	**52.1**	**75.6**	**70.9**

Table 5: Validation results sorted by increasing BLEU score. The best-performing setting (distilled NLLB-1.3B + fine-tuning) is shown last for emphasis.

5.5 Computational Requirements

NVIDIA T4 GPU (16GB) on Kaggle: Training time of about 5 hours (prematurely cut off at epoch 12). Memory 14.2GB with FP16. An inferred rate of about 167 sentences per minute. The test set of 5,000 sentences took approximately 6 minutes to complete. Model 2.6GB. Single T4 deployment was made possible with FP16. for faster training, we also used NVIDIA A100.

6 Experiments and Results

6.1 Model Comparison

Table 5 presents the validation results, showing that larger models achieve better performance. The 1.3B model achieves a BLEU score of +7 over the 600M, primarily due to its ability to handle complex legal patterns. Fine-tuning contributes approximately 4.0 BLEU points, which demonstrates its essential role in legal data adaptation. The distilled models are also performing. The `facebook/NLLB-1.3B-distilled` model achieves a score of 52.1 BLEU, compared to the standard version's score of 51. This is because distillation enhances generalization and alleviates overfitting. The general multilingual models are also more effective than the domain-specific models. The NLLB-1.3B-distilled model (52.1 BLEU) outperforms InLegalTrans (48.1 BLEU) due to its more extensive training and multilingual nature.

6.2 Ablation Studies

Table 6 quantifies design choices through ablation studies. The largest factor facilitating domain adaptation is fine-tuning (+4.0). The model size (+6.3) is worth the cost of computation to ensure legality. Complex legal sentences require long context (+1.9). Optimal decoding with beam search (+1.3). N-gram penalty +(0.8) does not allow repetition in legal formulae. Conservative LR (+1.4) maintains the knowledge of multilingualism. Distillation (+0.7) helps to improve performance by refining representations.

Configuration	BLEU	Δ
Full System	**52.1**	—
w/o Fine-tuning	39.2	-4.0
w/ NLLB-600M	43.2	-6.3
Max Length = 256	50.1	-1.9
Beam Width = 1	50.1	-1.3
No n-gram penalty	52.7	-0.8
LR = 2e-5	52.1	-1.4
Standard NLLB-1.3B	52.1	-0.7

Table 6: Ablation results. The distilled variant provides +0.7 BLEU improvement over the standard version.

7 Conclusion & Future Works

We presented Team-SVNIT's winning system for JUST-NLP 2025 Legal MT, achieving 1st place (AutoRank 61.62). Our approach demonstrates that carefully fine-tuned distilled multilingual models (`facebook/NLLB-1.3B-distilled`) outperform both smaller models and domain-specific systems with adequate training data (50K pairs) and systematic optimization. The system incorporates enhanced preprocessing to handle noisy legal texts and an optimized training setup with extended context (512), a conservative learning rate (2e-5), large batches (512), and cosine scheduling for stable convergence. Extensive ablation experiments quantify the impact of these design choices. Overall, the results demonstrate that domain adaptation through fine-tuning remains essential, and that large, well-pretrained multilingual models like NLLB-1.3B-distilled can outperform domain-specific models when sufficient fine-tuning data enables effective adaptation. Future research should address rare terminology through lexical constraints, code-switching through explicit guidelines, very long sentences through hierarchical approaches, multi-reference evaluation for accurate assessment, and document-level translation for improved consistency.

References

Kshetrimayum Boynao Singh, Deepak Kumar, and Asif Ekbal. 2025. Evaluation of LLM for English to Hindi Legal Domain Machine Translation Systems. In *Proceedings of the Tenth Conference on Machine Translation*, pages 823–833, Suzhou, China. Association for Computational Linguistics.

Marta R Costa-jussà, James Cross, Onur Çelebi, et al. 2022. No language left behind: Scaling human-centered translation. *arXiv:2207.04672*.

Marta R Costa-jussà, James Cross, et al. 2024. Scaling neural machine translation to 200 languages. *Nature*, 626:11–18.

Yeskendir Koishekenov, Alexandre Berard, and Vassilina Nikoulina. 2023. Memory-efficient NLLB-200 distilled models. In *Proc. ACL 2023*, pages 8031–8050.

Gowtham Ramesh, Sumanth Doddapaneni, Anoop Kunchukuttan, et al. 2021. Samanantar: Largest parallel corpora for Indic languages. *arXiv:2104.05596*.

Jay Gala, Pranjal A Chitale, Raghavan AK, et al. 2023. IndicTrans2: High-quality MT for 22 Indian languages. *arXiv:2305.16307*.

Argyri Panezi and John O'Shea. 2023. How can we manage the risks and liabilities associated with legal translation in the age of machine translation and generative AI? In *Generative AI and Law Workshop (ICML 2023)*.

Catherine Way. 2016. Challenges of legal translation in 21st century. *Intl. Journal of Communication*, 10:1009–1035.

Abdel Rahman Altakhaineh. 2025. Human versus AI translation of legal documents. *Intl. J. Language and Law*, 14(1):23–47.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: Automatic evaluation method for translation. In *Proc. ACL 2002*, pages 311–318.

Satanjeev Banerjee and Alon Lavie. 2005. METEOR: Automatic metric for evaluation. In *Proc. ACL Workshop*, pages 65–72.

Tianyi Zhang, Varsha Kishore, Ruoficer Luan, et al. 2019. BERTScore: Evaluating text with BERT. *arXiv:1904.09675*.

Ricardo Rei, Craig Stewart, Alon Lavie, et al. 2020. COMET: Neural framework for evaluation. In *Proc. EMNLP 2020*, pages 2685–2702.

Maja Popović. 2017. chrF++: Character n-gram F-score. In *Proc. WMT 2017*, pages 612–618.

Matthew Snover, Bonnie Dorr, Richard Schwartz, et al. 2006. Study of translation edit rate evaluation. In *Proc. AMTA 2006*, pages 223–231.

Ilya Loshchilov and Frank Hutter. 2017. SGDR: Gradient descent with restarts. In *Proc. ICLR 2017*.

Linjing Han, Qin Jin, and Michelle Yuan. 2024. Neural MT of clinical text documents. *BMC Medical Informatics and Decision Making*, 24:58.

Rahul Raja and Arpita Vats. 2025. Parallel corpora for low-resource languages. *arXiv:2503.04797*.

P Naveen, Rahul Aralikatte, Anoop Kunchukuttan, et al. 2024. Challenges of MT for low-resource languages. *iScience*, 27(9):110685.

Ashwini R Nair, Sukanta Roy, Swapnil Parekh, et al. 2024. Investigating translation for Indic languages. *Scientific Reports*, 14:17843.

Natalie Llop. 2025. AI, machine translation, and access to justice. *Stanford Law School Working Paper*.

Diptesh Suman, Vivek Srivastava, Swapnil Parekh, et al. 2023. Machine translation for low-resource languages. In *Proc. WMT 2023*, pages 856–864.

Shailashree K Sheshadri and KP Soman. 2023. A voyage on neural machine translation for Indic languages. *Procedia Computer Science*, 218:2694–2712.

Soham Bhattacharjee, Baban Gain, and Asif Ekbal. 2024. Domain dynamics: LLMs in English-Hindi translation. In *Proc. WMT 2024*, pages 341–354.

Markus Freitag, Nitika Mathur, Daniel Deutsch, et al. 2024. Are LLMs breaking MT metrics? WMT24 analysis. In *Proc. WMT 2024*, pages 47–81.

Ashish Vaswani, Noam Shazeer, Parmar Nikita, et al. 2017. Attention is all you need. In *Proc. NeurIPS 2017*, pages 5998–6008.

Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling knowledge in neural networks. *arXiv:1503.02531*.

Baban Gain, Ramakrishna Appicharla, Soumya Chennabasavaraj, et al. 2022. Low resource chat translation: Hindi-English benchmark. In *Proc. AMTA 2022*, pages 83–96.

Ramakrishna Appicharla, Asif Ekbal, and Pushpak Bhattacharyya. 2021. EduMT: MT system for educational content. In *Proc. ICON 2021*, pages 35–43.

Princeton NLP Group. 2025. Managing Context in Low-Resource Legal Machine Translation: A Case Study in English-Hindi. In *Proceedings of JUST-NLP 2025 Shared Task*.

Pratik Joshi, Sumanth Paul, Mukund Sharma, et al. 2024. IL-TUR: Benchmark for Indian legal understanding. *arXiv:2407.05399*.

Sayan Mahapatra, Debtanu Datta, Shubham Soni, et al. 2024. MILPaC: Benchmark for legal text translation. *ACM TALLIP*.

Ramakrishna Appicharla, Baban Gain, Santanu Pal, and Asif Ekbal. 2025. Beyond the sentence: Context-aware MT survey. *arXiv:2506.07583*.

Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, et al. 2024. Preliminary WMT24 ranking of MT systems. *arXiv:2407.19884*.

Appendix

A Result Analysis

A.1 Qualitative Analysis through Translation Examples

To better understand the performance of our model, we conducted a detailed manual analysis of the translation outputs. Table 7 presents representative examples showcasing both strengths and limitations.

The examples reveal several important patterns. Our model excels at translating standard legal terminology and complex sentence structures, particularly for common legal procedures and statutory references. The preservation of numerical data is consistently better. However, challenges remain with Latin legal terms (*quantum meruit*) and specialized legal concepts that require cultural adaptation rather than direct transliteration.

Nearly half of the translation mistakes occurred because the model couldn't handle unusual legal terms, particularly Latin phrases and specific legal concepts. This often means the system didn't "understand" a word, translated it incorrectly, or omitted it if it wasn't common in its training data. These problems arise frequently in legal writing, where precise terminology is crucial. To determine this percentage, we reviewed all mistakes in a batch of sample translations and counted the number that involved rare terminology. To reduce these errors, you'll need better resources or databases for legal terms, so the model knows what they mean. The next biggest issue (20%) came from the system mixing up pronouns in sentences with multiple people or actors, which shows it sometimes "loses track" of who is being referred to in complicated legal sentences.

A.2 Comparative Advantage of Distilled Models

Many reasons justify the superior performance of the distilled NLLB-1.3B variant, +0.7 BLEU over the standard variant(facebook/nllb-200-1.3B). First, knowledge distillation during pretraining definitely forces the model to learn more generalized representations rather than memorizing training patterns. Second, distilled models exhibit better calibration and reduced overconfidence, which is a crucial requirement for legal translation in accurately representing uncertainty. Third, the distillation process appears to enhance cross-lingual transfer efficiency, which is particularly beneficial in the case of English-Hindi legal translation due to the limited amount of available parallel data.

A.3 Practical Implications and Deployment Considerations

Our findings have significant practical consequences for legal translation workflows. The low TER score of 37.09 indicates that post-editing effort would be considerably reduced, allowing translator productivity to increase 2-3×. Numerical data and citations are perfectly preserved, which eliminates critical risks in legal documentation. The 40 to 45 % error rate confirms that human review is still important for legally binding documents.

The model size is moderate at 2.6GB, with computational demands for deployment in-house, which addresses all data confidentiality concerns usually associated with legal practice. The relatively short training time of under 3 hours enables organizations to fine-tune the model for their specific legal subdomains, such as patent law or corporate contracts.

English Source	**Hindi Translation**
the appellant is acquitted.	अपीलार्थी को बरी किया जाता है
they also raised memorials on the merits and the preliminary habit.	उन्होंने गुणावगुणों तथा प्रारंभिक अभ्यापत्ति पर भी सम्प–रीक्षण किये थे ।
being aggrieved by the order dated 2nd March , 2012 made by the learned Single Judge in CWJC No.3653 of 2012 , the writ petitioner has filed this appeal under clause 10 of the Letters Patent .	CWJC सं ० 3653 वर्ष 2012 में विद्वान एकल न्याया–धीश द्वारा किये गये दिनांक 2 मार्च, 2012 के आदेश से व्यथित होकर, रिट याची ने लेटर्स पेटेंट के खंड 10 के अधीन यह अपील दाखिल किया है ।
6- The opposition no.2 has filed his presence in this Court by filing the right in favour of his learned counsel , though he has not filed any counter affidavit .	6 – विपक्षी संख्या 2 ने अपने विद्वान अधिवक्ता के पक्ष में अधिकार दाखिल करके इस न्यायालय में अपनी उपस्थिति दर्ज करायी है, यद्यपि उसने कोई प्रति शपथ पत्र दाखिल नहीं किया है ।
7- We have heard the counsel for the learned Principal Additional Advocate General, Muzaffarpur Properties Private Limited, Smt. Shahida Hassan and the counsels of various dignitaries who have filed applications in these appeals both on facts as well as on law.	7 – हमने विद्वान प्रधान महाधिवक्ता , मुजफ्फरपुर गुण प्राईवेट लिमिटेड , श्रीमती शाधा हसन के अधिवक्ताओं को सुना है जिन्होंने दोनों तथ्यों तथा तथ्यों पर भी तथा विधि पर भी इन अपीलों में आवेदन दाखिल किये हैं ।
the aforesaid case was of the Central Excise Act and section 35H of the Central Excise Act provided that an appeal and reference should be made to the High Court within 180 days from the date of communication of the judgment of the order.	पूर्वोक्त मामला केन्द्रीय उत्पाद अधिनियम की केंद्रीय उत्पाद अधिनियम एवं धारा 35H का था यह प्रावधान करता है कि आदेश के निर्णय की संसूचना की तिथि से 180 दिनों के भीतर उच्च न्यायालय को एक अपील एवं निर्देश दिया जाना चाहिए ।

Table 7: English-Hindi translation examples demonstrating model performance.