

An Evaluation Framework for Legal Document Summarization

Ankan Mullick^{♣†} Abhilash Nandy^{♣♦†} Manav Nitin Kapadnis^{♣†}
 Sohan Patnaik[♣] R Raghav[♣] Roshni Kar[♣]

♣Indian Institute of Technology Kharagpur ♦ L3S Research Center, Leibniz Universität Hannover
 {ankanm, nandyabhilash}@kgpian.iitkgp.ac.in,
 {iammanavk, sohanpatnaik106, rraghav5600, roshnikar}@iitkgp.ac.in

Abstract

A law practitioner has to go through numerous lengthy legal case proceedings for their practices of various categories, such as land dispute, corruption, etc. Hence, it is important to summarize these documents, and ensure that summaries contain phrases with intent matching the category of the case. To the best of our knowledge, there is no evaluation metric that evaluates a summary based on its intent. We propose an automated intent-based summarization metric, which shows a better agreement with human evaluation as compared to other automated metrics like BLEU, ROUGE-L etc. in terms of human satisfaction. We also curate a dataset by annotating intent phrases in legal documents, and show a proof of concept as to how this system can be automated. Additionally, all the code and data to generate reproducible results is available on Github.

Keywords: Summarization, Evaluation Methodologies, Information Extraction, Legal Dataset

1. Introduction

Summarization could be extractive, where the summary has spans from the original text or abstractive, where the summary is generated using the original text. (Sai et al., 2020), (Chen et al., 2019a) list various metrics to evaluate summarization. BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE (Lin, 2004) etc. are context-free metrics, which work well for extractive summarization, while SMS (Clark et al., 2019), BERTScore (Zhang* et al., 2020) etc. are context-based metrics, which work well for abstractive summarization. Automatic summarization of legal documents (Bhattacharya et al., 2019) is required because - (1) Average length of an any Court Judgment is as high as 4,500 words (for example - Indian Supreme Court Judgments) (2) A law practitioner has to go through all contents of previous legal proceedings manually (3) Hiring Legal experts to summarize legal documents is expensive and very time consuming. However, while evaluating the quality of summaries, existing metrics fail when evaluating the amount of intent in the original text that is captured by the summary (intent here refers to the intention latent in a piece of text. e.g. (a) ‘Accused No. 1 Balwan Singh (appellant in Criminal Appeal No. 727 of 2015), on 22nd January, 2007, at evening time, was talking with the other accused regarding preparation to kill’ - in this sentence, the phrase ‘preparation to kill’ depicts the intent of Murder (b) ‘In the case in hand, robbed articles were found to be kept concealed at a place within knowledge of the applicant/accused No.1, and therefore, he is presumed to be one of the decoit involved in the decoity at the house of the first Aarti Palkar’ - in this sentence, the phrases ‘robbed articles were found to be kept concealed’ and ‘involved in the decoity’ de-

pic the intent of Robbery). To tackle this problem, we propose a novel evaluation metric that takes the help of intent phrases annotated using legal case documents, such that, the intent of these phrases matches with the category of the case. We use this proposed metric (and other metrics) to evaluate unsupervised summarization methods on legal documents and compare this with human evaluation. Another contribution is the curation of a dataset that consists of 101 legal documents spanning four categories of intents - Corruption, Land Dispute, Murder, and Robbery, along with a list of annotated intent phrases per document. Additionally, we come up with a framework, showing that the annotation of intent phrases and classification of documents into categories of intents can be automated. You can test our methods using this demo website¹.

2. Related Work

Unsupervised Summarization: Unsupervised Approaches ((Verma and Nidhi,),(Polsley et al., 2016),(Miller, 2019)) use semantic and analytical signals from the text to point out significant sentences for summarization. CaseSummarizer (Polsley et al., 2016) is specific to the legal domain that builds on existing methods to present an interface with scalable summary text, lists of entities and abbreviations, and a significance heat map of the text. BERT Extractive Summarizer (Miller, 2019) yields sentence embeddings by sending tokenized sentences to BERT (Devlin et al., 2019a) which are passed through hidden layers to get document level features. Finally, the summary prediction is compared to the ground truth.

Supervised Summarization: Supervised approaches ((See et al., 2017),(Saravanan et al., 2006),(Farzindar and Lapalme, 2004),(led,),(leg,)) take in documents

[†]Authors contributed equally

¹Demo website: <https://bit.ly/demoLREC2022>

and ground truth summaries, and use sentence features (e.g., facts of the case, background etc.) to filter good candidates for inclusion in summary. Graphical CRF Model (Saravanan et al., 2006) is trained using lexical and syntactic features to classify parts of the document into different categories such as ‘facts’, ‘arguments’ etc. Then, a K-mixture model is used to rank sentences, and a summary of the desired length is the output. LetSum (Farzindar and Lapalme, 2004) extracts important sentences by connecting the topical structure in the document and certainty of contentious themes of sentences in the judgment. Longformer Encoder Decoder (LED) (Beltagy et al., 2020) and (led,) supports long document generative sequence-to-sequence tasks, making it simple to process documents of thousands of tokens or longer.

But none of these works are judged against the intent of the case. Our proposed intent metrics for legal document summarization shows better relevance and human judgmental scores.

3. Proposed Evaluation Metric

We introduce an intent-based F1-Score and Human Score (related to Spearman Rank Correlation) metric for evaluation of a summary, referred to as ‘Intent Metric’ hereon. We report the average Intent Metric over all documents. Let us define ‘closePair’ as a pair of intent phrase and a sentence from the summary, such that, the intent phrase is contained in the sentence. The fraction of sentences in the summary that form a ‘closePair’ with atleast one intent phrase gives precision. Similarly, fraction of intent phrases that form a ‘closePair’ with atleast one sentence from the summary gives recall. Finally, Intent Metric is the F1 Score obtained from the precision and recall values. Given a document, the corresponding set P of M intent phrases and output summary O consisting of N sentences, a similarity score s_{ij} between i_{th} intent phrase (P_i) and j_{th} sentence in the summary (O_j) is 1 if P_i is a phrase contained in O_j and 0 otherwise, $\forall i \in \{1, 2, \dots, M\}, \forall j \in \{1, 2, \dots, N\}$. Mathematically,

$$s_{ij} = \begin{cases} 1, & \text{if } \exists k, P_i = O_j[k : k + \text{length}(P_i)] \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

P_{int} , R_{int} and $F1_{int}$ are calculated in the following manner -

$$P_{int} = \frac{\sum_{j=1}^N \mathbf{1} \left[\sum_{i=1}^M s_{ij} > 0 \right]}{N} \quad R_{int} = \frac{\sum_{i=1}^M \mathbf{1} \left[\sum_{j=1}^N s_{ij} > 0 \right]}{M} \quad (2) \quad (3)$$

$$F1_{int} = \frac{2 \cdot P_{int} \cdot R_{int}}{P_{int} + R_{int}} \quad (4)$$

Additionally, we also measure Precision, Recall, and F1 Score Metrics for evaluation of Slot, Intent, and Document Classification in Section 5.

4. Dataset Description

5000 legal documents are scraped from CommonLII² using ‘selenium’ python package. 101 documents belonging to the categories of Corruption, Murder, Land Dispute, and Robbery are randomly sampled from this larger set.

In case of Australian dataset (abbreviated as ‘AD’), we downloaded the Legal Case Reports Dataset³ from the UCI Machine Learning repository. The annotators then manually annotate randomly taken 59 relevant documents belonging to Corruption, Murder, Land Dispute, and Robbery categories.

Intent phrases are annotated for each document in the following manner -

- Initial filtering:** 2 annotators filter out sentences that convey an intent matching the category of the document at hand.
- Intent Phrase annotation** 2 other annotators then extract a span from each sentence, so as to exclude any details do not contribute to the intent (such as name of the person, date of incident etc.), and only include the words expressing corresponding intent. The resulting spans are the intent phrases. Overall Inter-annotator agreement (Cohen κ) is 0.79.

Table 1 shows the statistics of both the datasets, describing the number of documents, average length of documents, and intent phrases for each of the 4 intent categories. The documents on Robbery and Land Dispute are roughly longer than those on Murder and Corruption.⁴

| Category | No. of docs | | Avg. no. of words/doc | | Avg. no. of sentences/doc | | Avg. no. of words/intent phrase | |
|--------------|-------------|----|-----------------------|-------|---------------------------|-----|---------------------------------|----|
| | ID | AD | ID | AD | ID | AD | ID | AD |
| Corruption | 19 | 15 | 2542 | 4613 | 197 | 264 | 6 | 6 |
| Land Dispute | 27 | 14 | 2461 | 11508 | 196 | 579 | 5 | 6 |
| Murder | 32 | 15 | 1560 | 3008 | 149 | 183 | 6 | 5 |
| Robbery | 23 | 15 | 1907 | 7123 | 162 | 449 | 4 | 5 |

Table 1: Statistics for each category in both the datasets (ID - Indian-Data, AD - Australian-Data). The numbers are rounded to the nearest integer.

5. Experiments and Results

Competing baselines of Summarization Methods⁵:

The following summarization Methods (discussed in Section 2) are used in an unsupervised setting -

²<http://www.commonlii.org/resources/221.html>

³<https://archive.ics.uci.edu/ml/datasets/Legal+Case+Reports>

⁴We have used NLTK and Spacy for data pre-processing.

⁵We used Pytorch and Tensorflow for model implementation

1. **Graphical Model** (Saravanan et al., 2006) - Model trained on annotated data released in (Bhattacharya et al., 2019) is used for inference.
2. **LetSum** (Farzindar and Lapalme, 2004) - The process suggested in (Bhattacharya et al., 2019) is used for inference.
3. **Legal-Longformer Encoder Decoder (Legal-LED)** (leg,) - Longformer Encoder Decoder (Beltagy et al., 2020) trained on sec litigation releases (sec,) is used for inference.
4. **BERT Extractive Summarizer** (Miller, 2019)

Document Classification

| Model Name | Accuracy | | Macro F1 | |
|---------------------|-------------|-------------|-------------|-------------|
| | ID | AD | ID | AD |
| Logistic Regression | 0.62 | 0.50 | 0.38 | 0.47 |
| SVM | 0.62 | 0.50 | 0.38 | 0.42 |
| AdaBoost | 0.81 | 0.67 | 0.78 | 0.58 |
| BERT | 0.70 | 0.75 | 0.69 | 0.64 |
| RoBERTa | 0.75 | 0.67 | 0.70 | 0.60 |
| ALBERT | 0.70 | 0.67 | 0.69 | 0.60 |
| DeBERTa | 0.75 | 0.67 | 0.71 | 0.60 |
| LEGAL-BERT | 0.80 | 0.75 | 0.79 | 0.73 |
| LEGAL-RoBERTa | 0.67 | 0.75 | 0.65 | 0.64 |

Table 3: Results of Document Classification.

Recent developments show that, Transformer (Vaswani et al., 2017) based pre-trained language models like BERT (Devlin et al., 2019b), RoBERTa (Liu et al., 2019), ALBERT (Lan et al., 2020), and DeBERTa (He et al., 2021), have proven to be very successful in learning robust context-based representations of lexicons and applying these to achieve state of the art performance on a variety of downstream tasks such as document classification in our case.

We implemented different machine learning and transformer-based models mentioned in Table 3. Furthermore, We also tried domain-specific LEGAL-BERT (Chalkidis et al., 2020) and LEGAL-RoBERTa⁶ which were pre-trained on large scale legal corpora which in turn led to much better scores than their counterparts pre-trained on general corpora.

We observe from Table 3 that boosting algorithms such as AdaBoost (Freund and Schapire, 1999) and domain pre-trained transformer models such as LEGAL-BERT outperforms all the other models in terms of Accuracy and Macro F1-score in both the ID and AD datasets.

All of the transformer models were implemented using sliding window attention (Masood et al., 2020), since the document length for all the documents is greater than the transformer maximum token size. They were

⁶<https://huggingface.co/saibo/legal-roberta-base>

trained with a sliding window ratio of 20% over three epochs with learning rate and batch size set at 2×10^{-5} and 32 respectively. The documents in the dataset are randomly split into train, validation and test sets in the ratio of 6:2:2. The machine learning models were implemented on the TF-IDF features extracted from of the document text.

Intent Classification Using JointBERT

| Model Name | Accuracy | | Macro F1 | |
|-----------------|-------------|-------------|-------------|-------------|
| | ID | AD | ID | AD |
| JointBERT | 0.89 | 0.85 | 0.88 | 0.84 |
| JointDistilBERT | 0.95 | 0.70 | 0.95 | 0.69 |
| JointALBERT | 0.89 | 0.71 | 0.87 | 0.68 |

Table 3: Results on Intent classification.

We used Joint-BERT (Chen et al., 2019b) model on both the 'Indian-Data' as well as 'Australian-Data' for the task of intent classification between the classes of 'Corruption', 'Land Dispute', 'Robbery' and 'Murder'. The dataset is prepared in the following manner - Since there is a majority of sentences that have no intent phrase, only sentences containing an intent phrase, the one before that, and the one after that are used for training to mitigate class imbalance. Each sentence with an intent phrase has a target intent. The dataset is further randomly split into train (60%), validation (20%) and test sets (20%).

The different variations of JointBERT model perform reasonably well on the intent classification task for both the datasets, as seen from Table 3.

5.1. Evaluation using automated metrics

The following baseline metrics are used for comparison with the proposed metric -

BLEU (Papineni et al., 2002): It computes the number of n-grams in the predicted and reference summary. Overall score is found by taking the geometric mean of scores for n from 1 to 4.

METEOR (Banerjee and Lavie, 2005): It is F-measure based metric operating on unigrams by aligning and mapping each token in the predicted summary to a token in a reference summary.

ROUGE-L (Lin, 2004): It is an F-measure metric based on the longest common subsequence (LCS) between the reference and generated summary.

Sentence and Word Mover Similarity (S+WMS) (Clark et al., 2019): A linear programming solution measures the distance a predicted summary's embedding has to be moved to match the reference, and the similarity metric is calculated.

BERTScore (Zhang* et al., 2020): It obtains BERT (Devlin et al., 2019a) representations of each word in the predicted and reference summaries. Finally, a modified F1 score (weighted using inverse-document-frequency values) is found.

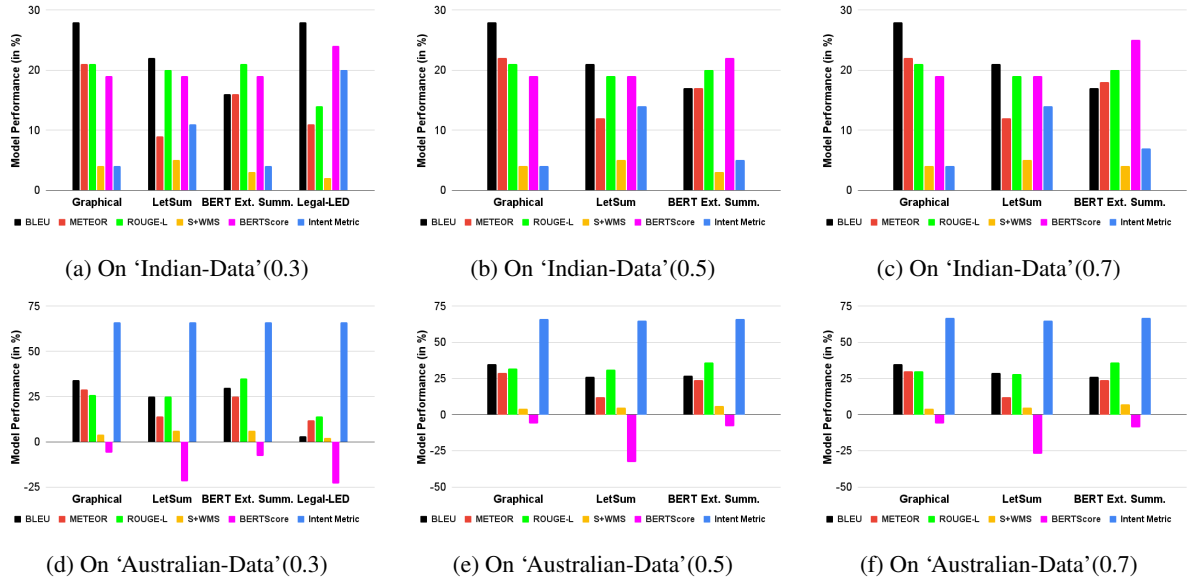


Figure 1: Model evaluation results on different ratios⁷ mentioned in parenthesis on ID and AD.

Fig. 1 plots the evaluation metrics for the two datasets and different lengths of summary as a fraction of the original document length (fractions are 0.3, 0.5, 0.7). In some cases, BERTScore is negative as BERTScore ranges from -1 to 1 . Also, *Legal-LED consumes more than 95% of GPU memory when the summary length is 50% and 70% of the original text, and hence, could not be reported*. The scores do not depend significantly on the summary length as a fraction of the input. However, we cannot conclude if one metric is better than the other, as every metric has its own way of quantifying the summary quality. Comparing the three models - (1) Graphical Model tends to perform the best for lexical metrics such as BLEU, METEOR, ROUGE-L. (2) BERT Extractive Summarizer gives the best BERTScore, as is expected. (3) Legal-LED performs better on ‘Indian Data’ compared to ‘Australian Data’. (4) In case of ‘Indian Data’, LetSum performs the best as per Intent Metric and S+WMS, while in case of ‘Australian Data’, all models perform almost equally well w.r.t these metrics. (5) Given a dataset, Intent Metric significantly varies across different summarization methods, which makes Intent Metric human-readable. To compare the quality of metrics, we see how well they correlate with human judgement in Section 5.2.

Also, from the correlation matrices among all the evaluation metrics corresponding to both datasets in Fig. 2, we find that Intent Metric has the highest correlation of with S+WMS in case of ‘Indian Data’, and with BERTScore in case of ‘Australian-Data’. Hence, our metric shows high correlation with metrics that quantify semantic similarity, rather than lexical similarity.

⁷Ratio here is the ratio of (length of summary/length of original text)

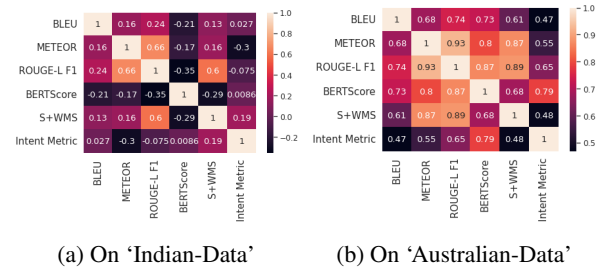


Figure 2: Correlation Matrix for Intent Metric and other baseline metrics

We perform our experiments on server with a RAM of 12.69 GB and a NVIDIA Tesla K80 GPU with a 12 GB memory.

5.2. Human Evaluation

To validate an automated evaluation metric, human evaluation of the generated summaries is necessary. We use Appen (<https://client.appen.com/>) platform to carry out the survey (<https://bit.ly/3n7xbCb>). As discussed in (Chen and Bansal, 2018), (Guo et al., 2018), measuring Relevance (if the summary contains salient information from original text) and Readability (coherence and fluency of the summary) of the summaries are essential for evaluating the quality of the summary. We report Relevance and ‘Human Score’, which is the average of Relevance and Readability.

For a survey on each dataset, 40 documents in case of ‘Indian Dataset’, and 20 documents in case of ‘Australian Dataset’ are sampled (each document has less than 20,000 characters to reduce annotation load). These documents are randomly split into 4 equal-sized sets, and for each set, a different summarization method is used. For each (original text, summary) pair, 3 ques-

| Model Name | BLEU | | METEOR | | ROUGE-L F1 | | BERT Score | | S+WMS | | Intent Metric | |
|-------------|-------|--------------|--------|-------------|------------|-------|------------|-------|-------|-------|---------------|-------|
| | ID | AD | ID | AD | ID | AD | ID | AD | ID | AD | ID | AD |
| Relevance | -0.09 | -0.03 | -0.14 | -0.09 | 0.06 | -0.32 | 0.03 | -0.18 | 0.25 | -0.59 | 0.42 | -0.05 |
| Human Score | -0.02 | 0.09 | -0.03 | 0.09 | 0.18 | -0.21 | -0.04 | 0.04 | 0.19 | -0.57 | 0.34 | -0.04 |

Table 5: Spearman Rank Correlation of automated metrics with human evaluation metrics on both ID ('Indian-Data') and AD ('Australian-Data'). Highest correlation corresponding to each dataset and human evaluation metric is in **bold**.

tions are asked - (1) category of the legal case (2) Relevance (3) Readability. For Relevance and Readability, the annotator has to pick from a 1 – 5 Likert Scale ('Very Poor' - 1, 'Poor' - 2, 'Fair' - 3, 'Good' - 4, 'Excellent' - 5). One document is annotated by two annotators. Average inter-annotator agreement Cohen κ is 0.74.

From Table 5, in case of 'Indian-Data', Intent Metric beats other metrics in both 'Relevance' as well as 'Human Score'. In case of the 'Australian-Data', the correlation of Intent Metric 'Relevance' and 'Human Score' is second and third best from the highest one in both fields. However, the average correlation across the two datasets is the highest among all metrics w.r.t both Relevance (0.185) and Human Score (0.15). Hence, we can conclude that Intent Metric is an important metric in terms of overall human satisfaction.

5.3. Working Demonstration

This section elaborates the process by which users can implement our methods.

Figure 3 shows the landing page of the demonstration website.

An Evaluation Framework for Legal Document Summarization

This demonstration can perform three different tasks:

- Summarize your document using 4 different models, namely:
 - Graphical Model (Saravanan et al., 2006)
 - LetSum (Farzindar et al., 2004)
 - BERT Extractive Summarizer (Devlin et al., 2018)
 - Legal-Longformer Encoder Decoder (Legal-LED) (Beltagy et al., 2020)
- Extraction of Intent from the uploaded documents using JointBERT (Chen et al., 2019)
- Evaluation of summary generated by one or more selected from the above models

Example Test File 1 : <https://drive.google.com/file/d/11siswn37SeeZJl6vnB3FVwqS2k0YIF/view>

Example Test File 2 : <https://drive.google.com/file/d/1S0QsZwXBIG78fA26OMDjkmnU7nDZYSm/view>

Upload the text (.txt) file that you would like to summarize:

Please upload a text(.txt) file (containing not more than 2000 words)

Figure 3: Landing Page of Demonstration

The demonstration can perform three different tasks:

- Summarize the document you upload with either one of the four available options:
 - Graphical Model** (Saravanan et al., 2006)

- LetSum** (Farzindar and Lapalme, 2004)
- Legal-Longformer Encoder Decoder (Legal-LED)** (leg,)
- BERT Extractive Summarizer** (Miller, 2019)

- Extraction of Intent from the uploaded document data using Joint-BERT (Chen et al., 2019b)
- Evaluation of summary generated by the chosen model.

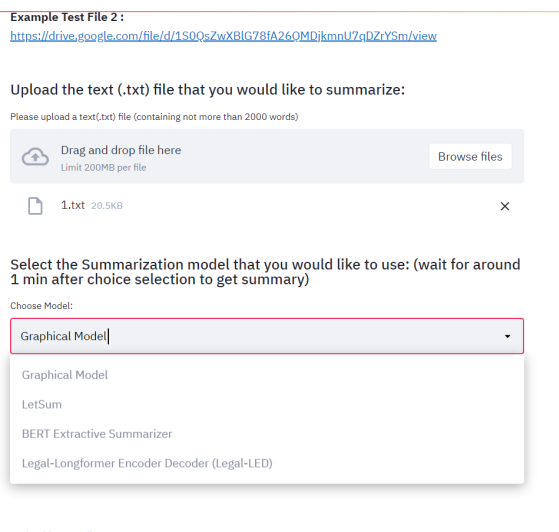


Figure 4: File upload and Model Selection Options

The user can upload their own text file, or can use any one of the two example files whose link are present on the demonstration page. Furthermore, after uploading the text file, the user has to select any one of the four options available in the dropdown list of Figure 4 and then select the "Click to start Summarization" option in order to run the model to start summarization of the uploaded text.

After the model is selected in Figure 4, the model is instantiated, and after sometime output summary is shown as output in the green box as seen in Figure 5.

Select the Summarization model that you would like to use: (wait for around 1 min after choice selection to get summary)

Choose Model:

LetSum

Click to start summarization

Running letsum(...)

Road, House State Police Station Station appellant accused convicted murder one Shailesh offence Section 302 appellant First appellant appellant sentence life imprisonment motive members appellant towards family deceased appellant son brought Rafiq stated assault appellant deceased. The appellant deceased appellant blow knife appellant suggest finding deceased reading evidence, find evidence Rafiq assault deceased appellant help us give benefit case case direct evidence, point motive may evidence point commission suicide Jitendra appellant deceased suicide motive appellant commit family deceased commission suicide son, us hold appellant deceased would turn intention knowledge Section 300 addition appellant also taking forehead may due fall stab wounds true appellant given also come evidence appellant ran 11 January, appellant brought knife find intended deceased postmortem prosecution prove Section 300 intention appellant inflict cause death Shailesh appellant intention deal injury cause death appellant knowledge injury would result death appellant intention injury may death case Section 300 going judgment trial Court, find considered say appellant intended inflict injury deceased reason stated given benefit doubt guilt Section 302 IPC find prosecution guilt appellant Section punishment Section 304 I incident due appellant time, age appellant years also may also fine Rs.5,25,000 may Civil

Figure 5: Summarization Output

The following are a set of extracted intent phrases by using JointBERT:

(Please wait for around 2 minutes since the model is being run on CPU)

- 1.- IN THE HIGH COURT OF JUDICATURE AT BOMBAY
- 2.- Aged about 60 years, Occu. Electrician,
- 3.- Near the House of Rajaram Admane, Nagpur.
- 4.- Shri R.M. Daga, Advocate for the Appellant.
- 5.- Smt. A.R. Kulkarni, A.P.P. for the Respondent/State.
- 6.- S.B. SHUKRE & S.M. MODAK, JJ.
- 7.- ORAL JUDGMENT (Per S.M. Modak, J.)
- 8.- The appellant sole accused is convicted of committing the murder of one Shailesh Balkrishna
- 9.- Junghare. He was assaulted with the help of knife on 23/09/2014 at about 7.15 p.m. in front of
- 10.- Nagpur initially registered an offence under Section 307 of the Indian Penal Code (for short
- 11.- hereinafter referred to as 'IPC'). Deceased Shailesh was first taken to the hospital of Dr. Patil and he
- 12.- died there lateron. Police converted the offence to Section 302 of IPC.
- 13.- respondent. They have taken us through the evidence. The emphasis of the appellant is entirely on

Figure 6: Extracted Intent Phrases from JointBERT

Once, the data is summarized using the model of user's choice. The demonstration automatically instantiates JointBERT (Chen et al., 2019b) for automated intent phrase extraction from the original data as seen in Figure 6 for further evaluation using our proposed Intent Metric. Furthermore, JointBERT also performs intent classification which gives the percentage of each intent present in the uploaded document whose output could be seen in Figure 7.

153.- directed to undergo rigorous imprisonment for two years.

154.- vi. Sentences to run concurrently.

155.- vii. Appellant be given benefit of set off. viii. Amount of fine be deposited in District

156.- Court, Nagpur. As and when deposited, compensation be paid to Reena Shailesh

The results of percentages of different intents are shown in the table:

| | Corruption | Land_Dispute | Murder | Robbery | UNK |
|---|------------|--------------|---------|---------|--------|
| 0 | 71.5596 | 16.0550 | 10.0917 | 1.3761 | 0.9174 |

The following are the reported metrics for the summary and the intent phrases extracted:

```
{
  "F1-score-rouge-I" : 0.07112375172682021
  "Precision-rouge-I" : 0.04655493482389125
  "Recall-rouge-I" : 0.15060240963855423
  "Bleu-score" : 0.363216360930019e-156
  "Meteor-score" : 0.0089863820547953387
  "Intent-F1-score" : 1
  "Intent-Recall" : 1
  "Intent-Precision" : 1
}
```

Figure 7: Intent Classification and Evaluation Results

6. Conclusion

In this paper, we explore a far less studied problem of devising a suitable evaluation metric for legal document summarization. To tackle the problem, a pre-condition is to curate a dataset that contains intent phrases extracted from legal documents belonging to categories like Robbery, Land Dispute, etc. This helps to develop a metric that correlates with human readability and relevance comparatively more than other metrics. We show a proof of concept that such intent phrase annotations required for the calculation of Intent Metric can be automated (Australian-Data). We believe that, such a metric would serve a better purpose in evaluating summarization of legal documents. We plan to extend the work on different categories of legal documents for various countries. We shall make the code, data available after acceptance.

References

- Banerjee, S. and Lavie, A. (2005). Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Beltagy, I., Peters, M. E., and Cohan, A. (2020). Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Bhattacharya, P., Hiware, K., Rajgaria, S., Pochhi, N., Ghosh, K., and Ghosh, S. (2019). A comparative study of summarization algorithms applied to legal case judgments. In *European Conference on Information Retrieval*, pages 413–428. Springer.
- Chalkidis, I., Fergadiotis, M., Malakasiotis, P., Aletas, N., and Androutsopoulos, I. (2020). Legal-bert: The muppets straight out of law school.
- Chen, Y.-C. and Bansal, M. (2018). Fast abstrac-

- tive summarization with reinforce-selected sentence rewriting. In *Proceedings of ACL*.
- Chen, A., Stanovsky, G., Singh, S., and Gardner, M. (2019a). Evaluating question answering evaluation. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 119–124, Hong Kong, China, November. Association for Computational Linguistics.
- Chen, Q., Zhuo, Z., and Wang, W. (2019b). Bert for joint intent classification and slot filling. *arXiv preprint arXiv:1902.10909*.
- Clark, E., Celikyilmaz, A., and Smith, N. A. (2019). Sentence mover’s similarity: Automatic evaluation for multi-sentence texts. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2748–2760.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019a). Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019b). Bert: Pre-training of deep bidirectional transformers for language understanding.
- Farzindar, A. and Lapalme, G. (2004). Letsum, a text summarization system in law field. In *a THE FACE OF TEXT conference (Computer Assisted Text Analysis in the Humanities)*, pages 27–36.
- Freund, Y. and Schapire, R. E. (1999). A short introduction to boosting.
- Guo, H., Pasunuru, R., and Bansal, M. (2018). Soft layer-specific multi-task summarization with entailment and question generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 687–697, Melbourne, Australia, July. Association for Computational Linguistics.
- He, P., Liu, X., Gao, J., and Chen, W. (2021). DeBERTa: Decoding-enhanced bert with disentangled attention.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. (2020). Albert: A lite bert for self-supervised learning of language representations.
-). Link to the led huggingface model. <https://huggingface.co/allenai/led-base-16384>.
-). Link to the legal led huggingface model. <https://huggingface.co/nsi319/legal-led-base-16384>.
- Lin, C.-Y. (2004). ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July. Association for Computational Linguistics.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach.
- Masood, M. A., Abbasi, R. A., and Wee Keong, N. (2020). Context-aware sliding window for sentiment classification. *IEEE Access*, 8:4870–4884.
- Miller, D. (2019). Leveraging bert for extractive text summarization on lectures. *arXiv preprint arXiv:1906.04165*.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Polsley, S., Jhunjhunwala, P., and Huang, R. (2016). CaseSummarizer: A system for automated summarization of legal texts. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: System Demonstrations*, pages 258–262, Osaka, Japan, December. The COLING 2016 Organizing Committee.
- Sai, A. B., Mohankumar, A. K., and Khapra, M. M. (2020). A survey of evaluation metrics used for NLG systems. *CoRR*, abs/2008.12009.
- Saravanan, M., Ravindran, B., and Raman, S. (2006). Improving legal document summarization using graphical models. *Frontiers in Artificial Intelligence and Applications*, 152:51.
-). Link to the sec litigation releases. <https://www.sec.gov/litigation/litreleases.htm>.
- See, A., Liu, P. J., and Manning, C. D. (2017). Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need.
- Verma, S. and Nidhi, V.). Extractive summarization using deep learning. *Advances in Artificial Intelligence*, page 107.
- Zhang*, T., Kishore*, V., Wu*, F., Weinberger, K. Q., and Artzi, Y. (2020). Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.