

BillSum: A Corpus for Automatic Summarization of US Legislation

Anastassia Kornilova

FiscalNote Research
Washington, DC

anastassia@fiscalnote.com

Vlad Eidelman

FiscalNote Research
Washington, DC

vlad@fiscalnote.com

Abstract

Automatic summarization methods have been studied on a variety of domains, including news and scientific articles. Yet, legislation has not previously been considered for this task, despite US Congress and state governments releasing tens of thousands of bills every year. In this paper, we introduce BillSum, the first dataset for summarization of US Congressional and California state bills (<https://github.com/FiscalNote/BillSum>). We explain the properties of the dataset that make it more challenging to process than other domains. Then, we benchmark extractive methods that consider neural sentence representations and traditional contextual features. Finally, we demonstrate that models built on Congressional bills can be used to summarize California bills, thus, showing that methods developed on this dataset can transfer to states without human-written summaries.

1 Introduction

The growing number of publicly available documents produced in the legal domain has led political scientists, legal scholars, politicians, lawyers, and citizens alike to increasingly adopt computational tools to discover and digest relevant information. In the US Congress, over 10,000 bills are introduced each year, with state legislatures introducing tens of thousands of additional bills. Individuals need to quickly process them, but these documents are often long and technical, making it difficult to identify the key details. While each US bill comes with a human-written summary from the Congressional Research Service (CRS),¹ similar summaries are not available in most state and local legislatures.

¹<http://www.loc.gov/crsinfo/>

Automatic summarization methods aim to condense an input document into a shorter text while retaining the salient information of the original. To encourage research into automatic legislative summarization, we introduce the BillSum dataset, which contains a primary corpus of 33,422 US Congressional bills and reference summaries split into a train and a test set. Since the motivation for this task is to apply models to new legislatures, the corpus contains an additional test set of 1237 California bills and reference summaries. We establish several benchmarks and show that there is ample room for new methods that are better suited to summarize technical legislative language.

2 Background

Research into automatic summarization has been conducted in a variety of domains, such as news articles (Hermann et al., 2015), emails (Nenkova and Bagga, 2004), scientific papers (Teufel and Moens, 2002; Collins et al., 2017), and court proceedings (Grover et al., 2004; Saravanan et al., 2008; Kim et al., 2013). The later area is most similar to BillSum in terms of subject matter. However, the studies in that area either apply traditional domain-agnostic techniques or take advantage of the unique structures that are consistently present in legal proceedings (e.g precedent, law, background).²

While automatic summarization methods have not been applied to legislative text, previous works have used the text to automatically predict bill passage and legislators’ voting behavior (Gerrish and Blei, 2011; Yano et al., 2012; Eidelman et al., 2018; Kornilova et al., 2018). However, these studies treated the document as a “bag-of-words” and did not consider the importance of individual

²Kanapala et al. (2017) provide a comprehensive overview of the works in legal summarization.

sentences. Recently, documents from state governments have been subject to syntactic parsing for knowledge graph construction (Kalouli et al., 2018) and textual similarity analysis (Linder et al., 2018). Yet, to the best of our knowledge, BillSum is the first corpus designed, specifically for summarization of legislation.

3 Data

The BillSum dataset consists of three parts: US training bills, US test bills and California test bills. The US bills were collected from the **Govinfo** service provided by the United States Government Publishing Office (GPO).³ Our corpus consists of bills from the 103rd-115th (1993-2016) sessions of Congress. The data was randomly split into 28,408 train bills and 5014 test bills. For California, bills from the 2015-2016 session were scraped directly from the legislature’s website;⁴ the summaries were written by their Legislative Counsel.

The BillSum corpus focuses on mid-length legislation from 5000 to 20,000 character in length. We chose to measure the text length in characters, instead of words or sentences, because the texts have complex structure that makes it difficult to consistently measure words. The range was chosen because on one side, short bills introduce minor changes and do not require summaries. While the CRS produces summaries for them, they often contain most of the text of the bill. On the other side, very long legislation is often composed of several large sections. The summarization problem thus becomes more akin in its formulation to multi-document summarization, a more challenging task that we leave to future work. The resulting corpus includes about 28% of all US bills, with a majority of the removed bills being shorter than 5000 characters.

For the summaries, we chose a 2000 character limit as 90% of summaries are of this length or shorter; the limit here is, also, set in characters to be consistent with our document length cut-offs. The distribution of both text and summary lengths is shown in Figure 1. Interestingly, there is little correlation between the bill and human summary length, with most summaries ranging from 1000 to 2000 characters.

For a closer comparison to other datasets, Table

³<https://github.com/unitedstates/congress>

⁴<http://leginfo.legislature.ca.gov>

1 provides statistics on the number of words in the texts, after we simplify the structure of the texts.

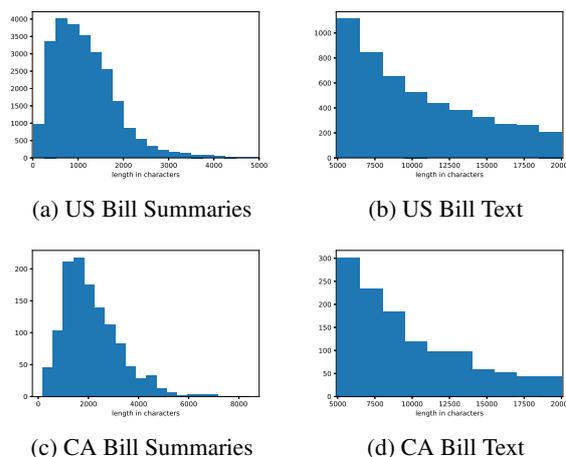


Figure 1: Bill Lengths

Stylistically, the BillSum dataset differs from other summarization corpora. Figure 2 presents an example Congressional bill. The nested, bulleted structure is common to most bills, where each bullet can represent a sentence or a phrase. Yet, content-wise, this is a straightforward example that states key details about the proposed grant in the outer bullets. In more challenging cases, the bill may state edits to an existing law, without whose context the change is hard to interpret, such as:

Section 4 of the Endangered Species Act of 1973 (16 U.S.C. 1533) is amended in subsection (a) in paragraph (1), by inserting “with the consent of the Governor of each State in which the endangered species or threatened species is present”

The average bill will contain both types of language, encouraging the study of both domain-specific and general summarization methods on this dataset.

4 Benchmark Methods

To establish benchmarks on summarization performance, we evaluate several extractive summarization approaches by first scoring individual sentences, then using a selection strategy to pick the best subset (Carbonell and Goldstein, 1998). While we briefly considered abstractive summarization models (Chopra et al., 2016), we found that the existing models trained on news and Wikipedia data produced ungrammatical results, and that the size of dataset is insufficient for the

		mean	min	25th	50th	75th	max
Words	US	1382	245	922	1254	1773	8463
	CA	1684	561	1123	1498	2113	3795
Sentences	US	46	7	31	41	58	372
	CA	47	12	31	42	59	137

Table 1: Text length distributions on preprocessed texts.

Summary. This bill authorizes the Department of Education to award competitive grants to nonprofit organizations for the development and implementation of teacher-led projects to improve outcomes in elementary and secondary schools. Grantee organizations shall use grant funds to make competitive subgrants to teachers and school leaders in partnership with the organization or a local educational agency.

SECTION 1. SHORT TITLE.

This Act may be cited as the "Teach to Lead Act of 2016".

SEC. 2. FINDINGS.

Congress finds as follows:

- (1) Teachers, because of their position in the classroom, often see important opportunities to improve student learning most directly and thus have a unique perspective from which to create practical solutions to help students succeed.
- (2) According to a Scholastic and Bill & Melinda Gates Foundation poll, 69 percent of teachers feel that their voices are heard in their school, but only one-third feel heard in their district, five percent in their State, and two percent at the national level.

SEC. 3. PURPOSE.

The purpose of this Act is to empower teachers to develop and implement projects with the potential to have a wider impact on developing the knowledge, pedagogical skills, and conditions needed to improve teaching and student outcomes, particularly academic growth, by bringing their classroom knowledge and expertise directly to bear on the many challenges confronting our education system.

SEC. 4. GRANT PROGRAM.

(a) In General.—

(1) **PROGRAM AUTHORIZED.**—From the funds made available under section 7, the Secretary of Education may make grants, on a competitive basis, to one or more nonprofit organizations to award subgrants to eligible entities to develop and implement teacher-led projects to improve teaching and student outcomes in elementary school and secondary school, particularly academic growth.

(2) **GRANT PERIOD.**—A grant made to a nonprofit organization under paragraph (1) shall be for a period of not more than five years.

(3) **USE OF GRANT FUNDS.**—A nonprofit organization that receives a grant under paragraph (1)—

- (A) shall reserve not less than 90 percent of the grant to award subgrants, on a competitive basis, to eligible entities under subsection (c); and
- (B) may use not more than 10 percent of the grant for administrative purposes.

(b) **Applications.**—A nonprofit organization that desires a grant under this section shall submit an application to the Secretary at such time and in such manner, and containing such information as the Secretary may require. The application shall—

- (1) demonstrate the entity's ability to—
 - (A) operate a national program, a multi-State program, or a program that reaches not less than 100,000 students;
 - (B) manage the administrative and fiscal aspects of the subgrant program described in this section; ...

(c) **Subgrants.**—

(1) **SUBGRANT PRIORITY.**—A nonprofit organization receiving a grant under this section shall use such grant to award subgrants to eligible entities under this subsection, and in awarding such subgrants the nonprofit organization shall give priority to eligible entities that will use the subgrants to carry out projects that—

- (A) are designed to improve teaching and learning outcomes for all students in high-need schools or that target the educational needs of low-income or minority students;

(2) **SUBGRANT APPLICATIONS.**—An eligible entity that desires a subgrant under this section shall submit an application to the applicable nonprofit organization awarded a grant under this section at such time and in such manner, and containing such information as the nonprofit organization may reasonably require. Each application shall, at a minimum, describe—

- (A) the project proposed, including timelines, resources needed, and any measurable objectives to be used in determining how the project will improve teaching and student outcomes, particularly academic growth...

(3) **USE OF SUBGRANT FUNDS.**—

(A) **USE OF SUBGRANT FUNDS.**—An eligible entity shall use the subgrant received under this section to develop and implement an innovative project designed and led by teachers, teams of teachers, or teachers and school leaders to improve teaching and learning at the elementary school and secondary school level, such as—

- (i) increasing student engagement through personalized learning, including technology-enabled instruction;
- (ii) strengthening support for educators, including support for implementation of challenging, academic standards to prepare students to be ready for college and careers...

(B) **ADMINISTRATIVE EXPENSES.**—A partner local educational agency or nonprofit organization that serves as the fiscal agent for an eligible entity, may use not more than two percent of the subgrant for direct administrative expenses incurred in carrying out its responsibilities under the subgrant.

SEC. 5. PERFORMANCE MEASUREMENT.

The Secretary shall establish goals and performance indicators to measure and assess the impact of the activities carried out under this Act.

SEC. 6. DEFINITIONS.

In this Act:

(1) **ELIGIBLE ENTITY.**—The term "eligible entity" means an individual teacher, a team of teachers, or teachers and school leaders, in partnership with a local educational agency or a nonprofit organization that serves as the fiscal agent with respect to funds awarded under this Act.

(2) **ESEA TERMS.**—The terms "elementary school", "secondary school", "local educational agency", and "Secretary" have the meanings given the terms in section 8101 of the Elementary and Secondary Education Act of 1965 (20 U.S.C. 7801).

SEC. 7. AUTHORIZATION OF APPROPRIATIONS.

There are authorized to be appropriated \$10,000,000 for each of the fiscal years 2017 through 2021 to carry out this Act.

Figure 2: Example US Bill

necessary retraining. Recent works have successfully fine-tuned models for other NLP tasks to specific domains (Lee et al., 2019), but we leave to future work the exploration of similar abstractive strategies.

The scoring task is framed as a supervised learning problem. First, we create a binary label for each sentence indicating whether it belongs in the summary (Gillick et al., 2008).⁵ We compute

⁵As noted in Section 3, it is difficult to define sentence boundaries for this task due to the bulleted structure of the documents. We simplify the text with the following heuristic: if a bullet is shorter than 10 words, we treat it as a part of the previous sentence; otherwise, we treat it as a full sentence. This cut-off was chosen by manually analyzing a sample of sentences. A more sophisticated strategy would be to check if each bullet is a sentence fragment with a syntactic parser and

a Rouge-2 Precision score of a sentence relative to the reference summary and simplify it to a binary value based on whether it is above or below 0.1 (Lin, 2004; Zopf et al., 2018). As an example, the sentences in the positive class are highlighted in green in Figure 2.

Second, we build several models to predict the label. For the models, we consider two aspects of a sentence: its importance in the context of the document (4.1) and its general summary-like properties (4.2).

4.1 Document Context Model (DOC)

A good summary sentence contains the main ideas mentioned in the document. Thus, researchers have designed a multitude of features to capture this property. We evaluate how several common ones transfer to our task:

The position of the sentence can determine how informative the sentence is (Seki, 2002). We encode this feature as a fraction of 'sentence position / total sentence count', to restrict this feature to the 0–1 range regardless of the particular document's length. In addition, we include a binary feature for whether the sentence is near a section header.

An informative sentences will contain words that are important to a given document relative to others. Following a large percentage of previous works, we capture this property using TF-IDF (Seki, 2002; Ramos et al., 2003). First, we calculate a document-level TF-IDF weight for each word, then take the average and the maximum of these weights for a sentence as features. To relate language between sentences, "sentence-level" TF-IDF features are created using each sentence as a document for the background corpus; the average and max of the sentence's word weights are used as features.

We train a random forest ensemble model over these features with 50 estimators (Breiman, 2001).⁶ This method was chosen because it best then reconstruct full sentences; however, the former approach is sufficient for most documents.

⁶Implemented with `scikit-learn.org`

captured the interactions between the small number of features.

4.2 Summary Language Model (SUM)

We hypothesize that certain language is more common in summaries than in bill texts. Specifically, that summaries primarily contain general effects of the bill (e.g awarding a grant) while language detailing the administrative changes will only appear in the text (e.g inserting or modifying relatively minor language to an existing statute). Thus, a good summary should contain only the major actions.

Hong and Nenkova (2014) quantify this aspect using hand-engineered features based on the the likelihood of words appearing in summaries as opposed to the text. Later, Cao et al. (2015) built a Convolutional Neural Network (CNN) to predict if a sentence belongs in the summary and showed that this straightforward network outperforms engineered features. We follow their approach, using the BERT model as our classifier (Devlin et al., 2018). BERT can be adapted for and has achieved state-of-the-art performance on a number of NLP tasks, including binary sentiment classification.⁷

To adapt the model to our domain, we pre-train the **Bert-Large Uncased** model on the “next-sentence prediction” task using the US training dataset for 20,000 steps with a batch size of 32.⁸ The pretraining strategy has been successfully applied to tune BERT for tasks in the biomedical domain (Lee et al., 2019). Using the pretrained model, the classification setup for BERT is trained on sentences and binary labels for 3 epochs over the training data.

4.3 Ensemble and Sentence Selection

To combine the signals from the DOC and SUM models, we create an ensemble averaging the two probability outputs.⁹

To create the final summary, we apply the Maximal Marginal Relevance (MMR) algorithm (Goldstein et al., 2000). MMR iteratively constructs a summary by including the highest scoring sentence with the following formula:

⁷All code described are used directly from <https://github.com/google-research/bert>

⁸This is the pretraining procedure recommended by the authors of BERT on their github website.

⁹Additional experiments using Linear Regression with the actual Rouge-2 Precision score as the target, but found that they produced similar results.

$$s_{next} = \max_{s \in D - S_{cur}} 0.7 * f(s) - 0.3 * sim(s, S_{cur})$$

where D is the set of all the sentences in the document, S_{cur} are the sentences in the summary so far, $f(s)$ is the sentence score from the model, sim is the cosine similarity of the sentence to S_{cur} , and 0.7 and 0.3 are constants chosen experimentally to balance the two properties. This method allows us to pick relevant sentences while minimizing redundancies. We repeat this process until we reach the length limit of 2000 characters.

5 Results

To estimate the upper bound on our approach, an oracle summarizer is created by using the true Rouge-2 Precision scores with the MMR selection strategy. In addition, we evaluate the following unsupervised baselines: **SumBasic** (Nenkova and Vanderwende, 2005), Latent Semantic Analysis (LSA) (Gong and Liu, 2001) and **TextRank** (Mihalcea and Tarau, 2004). The final results are shown in Table 2. The Rouge F-Score is used because it considers both the completeness and conciseness of the summary method.^{10,11}

We evaluated the DOC, SUM, and ensemble classifiers separately. All three of our models outperform the other baselines, demonstrating that there is a “summary-like” signal in the language across bills. The SUM model outperforms the DOC model showing that a strong language model can capture general summary-like features; this result is in line with Cao et al. (2015) and Collins et al. (2017) sentence level neural network performance. However, in those studies incorporating several contextual features improved the performance, while DOC+SUM performs similarly to DOC. In future work we plan to incorporate contextual features into the neural network directly; Collins et al. (2017) showed that this strategy is effective for scientific article summarization. In addition, we plan to explore additional sentence selection strategies instead of always adding sentences to the 2000 character limit.

Next, we applied our US model to CA bills. Overall, the performance is lower than on US bills

¹⁰Precision and recall scores are listed in the supplemental material for additional context.

¹¹Rouge scores calculated using <https://github.com/pcyin/PyRouge>

(Table 2b), but all three supervised methods perform better than the unsupervised baselines, suggesting that models built using the language of US Bills can transfer to other states. Interestingly, DOC+SUM outperforms DOC in the CA dataset, suggesting that the BERT model may have overfit to the US language. An additional reason for the drop in the performance is the difference in the structure of the summaries: In California the provided summaries state not only the proposed changes, but the relevant pieces of the existing law, as well (see Appendix B.3 for a more in-depth discussion). We hypothesize that a model trained on multi-state data would transfer better, thus we plan to expand the dataset to include all twenty-three states with human-written summaries.

Table 2: ROUGE F-scores (%) of different methods.

	Rouge-1	Rouge-2	Rouge-L
Oracle	47.97	32.14	40.89
SumBasic	30.56	15.33	23.75
LSA	32.24	14.02	23.75
TextRank	34.10	17.45	27.57
DOC	38.18	21.22	31.02
SUM	41.29	24.47	34.07
DOC + SUM	41.28	24.31	34.15

(a) Congressional Bills

	Rouge-1	Rouge-2	Rouge-L
Oracle	51.24	34.86	45.91
SumBasic	35.47	16.16	30.10
LSA	35.05	16.34	30.10
TextRank	35.81	18.10	30.10
DOC	37.32	18.72	31.87
SUM	38.67	20.59	33.11
DOC + SUM	39.26	21.16	33.77

(b) CA Bills

5.1 Summary Language Analysis

The success of the SUM model suggests that certain language is more summary-like. Following a study by Hong and Nenkova (2014) on news summarization, we apply KL-divergence based metrics to quantify which words were more summary-like. The metrics are calculated by:

1. Calculate the probability of unigrams appearing in the bill text and in the summaries ($P_t(w)$ and $P_s(w)$ respectively).
2. Calculate KL scores as : $KL_w(S|T) = P_s(w) * \ln \frac{P_s(w)}{P_t(w)}$ and the opposite.

A large value of $KL(S|T)$ indicates that the word is summary-like and $KL(T|S)$ indicates a text-like word. Table 3 shows the most summary-like and text-like words in bills and resolutions. For both document types, the summary-like words tend to be verbs or department names; the text-like words mostly refer to types of edits or background content (e.g “reporting the rise of.”). This follows our intuition about summaries being more action driven. While a complex model, like BERT, may capture these signals internally; understanding the significant language explicitly is important both for interpretability and for guiding future models.

Table 3: Examples of summary and text like words

Summary-like	prohibit, DOD, VA, allow, penalty, prohibit, EPA, eliminate, implement, require
Text-like	estimate, average, report, rise, section, finish, percent, debate

6 Conclusion

In this paper, we introduced BillSum, the first corpus for legislative summarization. This is a challenging summarization dataset due to the technical nature and complex structure of the bills. We have established several baselines and demonstrated that there is a large gap in performance relative to the oracle, showing that the problem has ample room for further development. We have also shown that summarization methods trained on US Bills transfer to California bills - thus, the summarization methods developed on this dataset could be used for legislatures without human written summaries.

References

- Leo Breiman. 2001. [Random forests](#). *Mach. Learn.*, 45(1):5–32.
- Ziqiang Cao, Furu Wei, Sujian Li, Wenjie Li, Ming Zhou, and WANG Houfeng. 2015. Learning summary prior representation for extractive summarization. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, volume 2, pages 829–833.
- Jaime Carbonell and Jade Goldstein. 1998. [The use of mmr, diversity-based reranking for reordering docu-](#)

- ments and producing summaries. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '98, pages 335–336, New York, NY, USA. ACM.
- Sumit Chopra, Michael Auli, and Alexander M. Rush. 2016. [Abstractive sentence summarization with attentive recurrent neural networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 93–98. Association for Computational Linguistics.
- Ed Collins, Isabelle Augenstein, and Sebastian Riedel. 2017. A supervised approach to extractive summarisation of scientific papers. *arXiv preprint arXiv:1706.03946*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Vlad Eidelman, Anastassia Kornilova, and Daniel Argyle. 2018. How predictable is your state? leveraging lexical and contextual information for predicting legislative floor action at the state level. *arXiv preprint arXiv:1806.05284*.
- Sean Gerrish and David M Blei. 2011. Predicting legislative roll calls from text. In *Proceedings of the 28th international conference on machine learning (icml-11)*, pages 489–496.
- Daniel Gillick, Benoit Favre, and Dilek Hakkani-Tür. 2008. The icsi summarization system at tac 2008. In *TAC*.
- Jade Goldstein, Vibhu Mittal, Jaime Carbonell, and Mark Kantrowitz. 2000. [Multi-document summarization by sentence extraction](#). In *Proceedings of the 2000 NAACL-ANLP Workshop on Automatic Summarization - Volume 4*, NAACL-ANLP-AutoSum '00, pages 40–48, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Yihong Gong and Xin Liu. 2001. Generic text summarization using relevance measure and latent semantic analysis. In *Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 19–25. ACM.
- Claire Grover, Ben Hachey, and Ian Hughson. 2004. [The HOLJ corpus. supporting summarisation of legal texts](#). In *COLING 2004 5th International Workshop on Linguistically Interpreted Corpora*, pages 47–54, Geneva, Switzerland. COLING.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, pages 1693–1701.
- Kai Hong and Ani Nenkova. 2014. Improving the estimation of word importance for news multi-document summarization. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 712–721.
- Aikaterini-Lida Kalouli, Leo Vrana, Vigile Marie Fabella, Luna Bellani, and Annette Hautli-Janisz. 2018. Cousbi: A structured and visualized legal corpus of us state bills. In *LREC 2018*, pages 7–14.
- Ambedkar Kanapala, Sukomal Pal, and Rajendra Pamula. 2017. Text summarization from legal documents: a survey. *Artificial Intelligence Review*, pages 1–32.
- Mi-Young Kim, Ying Xu, and Randy Goebel. 2013. Summarization of legal texts with high cohesion and automatic compression rate. In *New Frontiers in Artificial Intelligence*, pages 190–204, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Anastassia Kornilova, Daniel Argyle, and Vlad Eidelman. 2018. Party matters: Enhancing legislative embeddings with author attributes for vote prediction. *arXiv preprint arXiv:1805.08182*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *CoRR*, abs/1901.08746.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*.
- Fridolin Linder, Bruce A. Desmarais, Matthew Burgess, and Eugenia Giraudy. 2018. Text as policy: Measuring policy similarity through bill text reuse. *SSRN*.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*.
- Ani Nenkova and Amit Bagga. 2004. Facilitating email thread access by extractive summary generation. *AMSTERDAM STUDIES IN THE THEORY AND HISTORY OF LINGUISTIC SCIENCE*, page 287.
- Ani Nenkova and Lucy Vanderwende. 2005. The impact of frequency on summarization. *Microsoft Research, Redmond, Washington, Tech. Rep. MSR-TR-2005*, 101.
- Juan Ramos et al. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, volume 242, pages 133–142.

Murali Saravanan, Balaraman Ravindran, Assistant Professor, and Dr Raman. 2008. Automatic identification of rhetorical roles using conditional random fields for legal document summarization. *J. Artif. Intell. Law*.

Yohei Seki. 2002. Sentence extraction by tf/idf and position weighting from newspaper articles.

Simone Teufel and Marc Moens. 2002. Summarizing scientific articles: experiments with relevance and rhetorical status. *Computational linguistics*, 28(4):409–445.

Tae Yano, Noah A Smith, and John D Wilkerson. 2012. Textual predictors of bill survival in congressional committees. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 793–802. Association for Computational Linguistics.

Markus Zopf, Eneldo Loza Mencía, and Johannes Fürnkranz. 2018. Which scores to predict in sentence regression for text summarization? In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1782–1791. Association for Computational Linguistics.

A Additional ROUGE Scores

As discussed in the Results section, F-Scores encourage a balance between comprehensiveness and conciseness. However, as it is useful to analyze the precision and recall scores separately, both are presented in Table 4 for US Bills and in Table 5 for CA Bills. All tested methods favor recall, since they consistently generate a 2000 character summary, instead of stopping early when a concise summary may be sufficient. For both datasets, the difference in Recall between the Oracle and DOC+SUM summarizer is a lot smaller than for Precision; which suggests that a lot of useful summary content can be found with an extractive method. In future work, we will focus on extracting more granular snippets to improve precision.

B Additional Bill Examples

We highlight several example bills to showcase the different types of bills found in the dataset.

B.1 Complex Structure Example

In the Data section, we discussed some of the challenges with processing bills: complex formatting and technical language. Figure 3 is an excerpt from a particularly difficult example:

Table 4: ROUGE Scores of Congressional Bills

	Rouge-1	Rouge-2	Rouge-L
Oracle	41.42	27.77	39.41
SumBasic	24.35	11.54	22.15
LSA	27.00	12.99	24.70
TextRank	29.44	14.99	26.65
DOC	31.80	17.50	29.34
SUM	34.79	20.45	32.32
DOC + SUM	35.15	20.42	32.65

(a) Precision Scores

	Rouge-1	Rouge-2	Rouge-L
Oracle	65.98	46.18	62.38
SumBasic	47.55	21.97	43.14
LSA	46.13	22.61	43.05
TextRank	46.64	25.36	42.11
DOC	54.95	32.80	50.67
SUM	58.03	36.79	53.87
DOC + SUM	57.35	36.39	53.22

(b) Recall Scores

Table 5: ROUGE Scores of California Bills

	Rouge-1	Rouge-2	Rouge-L
Oracle	41.42	27.77	39.41
SumBasic	33.72	16.30	30.41
LSA	34.84	17.24	31.69
TextRank	36.66	19.28	32.61
DOC	36.44	19.10	33.14
SUM	38.15	21.23	34.52
DOC + SUM	39.14	22.07	35.87

(a) Precision Scores

	Rouge-1	Rouge-2	Rouge-L
Oracle	65.98	46.18	62.38
SumBasic	40.47	17.89	36.36
LSA	38.23	17.28	34.61
TextRank	37.79	18.97	33.50
DOC	41.17	20.53	37.19
SUM	42.21	22.20	38.04
DOC + SUM	42.17	22.59	38.24

(b) Recall Scores

- The text interleaves several layers of bullets. Lines 3, 15, 27 represent the same level (points (3) and (4) omitted for space); lines 16, 17, 19 and 21 go together, as well. These multiple levels need to be handled carefully,

1	SEC. 4. WOMEN'S BUSINESS CENTER PROGRAM.
2	(a) Women's Business Center Financial Assistance.—Section 29 of the Small Business Act (15 U.S.C. 656) is amended
3	(1) in subsection (a)—
4	(A) by striking paragraph (4);
5	(B) by redesignating paragraphs (2) and (3) as paragraphs (3) and (4), respectively;
6	(C) by inserting after paragraph (1) the following:
7	“(2) the term ‘eligible entity’ means—
8	“(A) a private nonprofit organization;
9	“(B) a State, regional, or local economic development organization;
10	“(C) a development, credit, or finance corporation chartered by a State;
11	“(D) a junior or community college, as defined in section 312(f) of the Higher Education Act of 1965; or
12	“(E) any combination of entities listed in subparagraphs (A) through (D);”;
13	(D) by adding at the end the following:
14	“(5) the term ‘women’s business center’ means a project conducted by an eligible entity under this section.”;
15	(2) in subsection (b)—
16	(A) by redesignating paragraphs (1), (2), and (3) as subparagraphs (A), (B), and (C), respectively..
17	(B) by striking “The Administration” and all that follows through “5-year projects” and inserting the following:
18	“(1) IN GENERAL, The Administration may provide financial assistance to an eligible entity to conduct a project;
19	(C) by striking “The projects shall” and inserting the following:
20	“(2) USE OF FUNDS.—The project shall be designed to provide training and counseling that meets the needs of women, especially socially or economically disadvantaged women, and shall”; and
21	(D) by adding at the end the following:
22	“(3) AMOUNT OF FINANCIAL ASSISTANCE.—
23	“(A) IN GENERAL.—Except as provided in subparagraph (B), the amount of financial assistance provided under this subsection to an eligible entity per project year shall be not more than \$250,000.
24	“(B) ADDITIONAL FINANCIAL ASSISTANCE.—
25	“(G) IN GENERAL.—The Administrator may award financial assistance under this subsection to an eligible entity in an amount that is more than \$250,000 in a given project year if the Administrator determines that the eligible entity
26	“(I) obtained more than \$250,000 in non-Federal contributions for that project year in accordance with subsection (c);
27	(5) by striking subsection (f) and inserting the following:
28	“(f) Applications And Criteria For Initial Financial Assistance.—
29	“(1) APPLICATION.—Each eligible entity desiring financial assistance under subsection (b) shall submit to the Administrator an application that contains—
30	“(A) a certification that the eligible entity—
31	“(i) has designated an executive director or program manager, who may be compensated using financial assistance under subsection (b) or other sources, to manage the women’s business center for which assistance under subsection (b) is sought.

Figure 3: US H.R.1680 (115th)

or the summarizer will extract snippets that can not be interpreted without context.

- Lines 22-26 both introduce new language for the law and use the bulleted structure.
- Line 27 states that the existing “subsection (f)” is being removed and replaced. While lines 28 onward state the new text, the meaning of the change relative to the current text is not clear.

The human-written summary for this bill was:

(Sec. 4) “Women’s business center” shall mean a project conducted by any of the following eligible entities:

- *a private nonprofit organization;*
- *a state, regional, or local economic development organization;*
- *a state-chartered development, credit, or finance corporation;*
- *a junior or community college; or*
- *any combination of these entities.*

The SBA may award up to \$250,000 of financial assistance to eligible entities per project year to

conduct projects designed to provide training and counseling meeting the needs of women, especially socially and economically disadvantaged women.

Most of the relevant details are capture in the text between lines 8-14 and 20-24. For examples similar to this one, the summary language is extracted almost directly from the text, but, parsing them correctly from the original structure is a non-trivial task.

B.2 Paraphrase Example

For a subset of the bills, the CRS will paraphrase the technical language. In these cases, extractive summarization methods are particularly limited. Consider the example in Figure 4 and its summary:

This bill amends the Endangered Species Act of 1973 to revise the process by which the Department of the Interior or the Department of Commerce, as appropriate, reviews petitions to list a species on the endangered or threatened species list. Specifically, the bill establishes a process for the appropriate department to declare a petition backlog and discharge the petitions when there is a backlog.

SEC. 2. DEFINITIONS.
Section 3 of the Endangered Species Act of 1973 (16 U.S.C. 1532) is amended...
(3) by adding at the end the following:
“(b) Definitions Related To Petitions.—In this Act:
“(3) BACKLOG SCHEDULE.—The term ‘backlog schedule’ means a comprehensive, regularly updated compendium of petitioned-for species that are the subject of a 90-day petition backlog or a 12-month petition backlog—
“(A) that consists of—
“(i) a list of petitions to add a species to a list of species under section 4(c), including petitions to move a species from the list of threatened species to the list of endangered species; and
“(ii) a list of petitions to remove a species from a list of species under section 4(c), including petitions to move a species from the list of endangered species to the list of threatened species; and
“(B) in which the petitions in each such list appear in the order in which the petitions were submitted to the Secretary.
“(4) BACKLOG PROCEDURES.—The term ‘backlog procedures’ means the actions taken by the Secretary—
“(A) under section 4(b)(3)(G) following the declaration of a 90-day petition backlog; or
“(B) under section 4(b)(3)(H) following the declaration of a 12-month petition backlog.
SEC. 3. BACKLOG DECLARATION AND PROCEDURES.
(a) In General.—Section 4(b)(3) of the Endangered Species Act of 1973 is amended by adding at the end the following:
“(f) The Secretary shall
“(1) declare a 90-day petition backlog at any time the total number of species for which a petition is presented to the Secretary under subparagraph (A) that has not been the subject of a finding by the Secretary within the timeframe established under such subparagraph exceeds 5 percent of the number of species for which such petitions have been presented during the preceding 15 years;
“(1) submit a backlog schedule for such backlog to—
“(aa) the President;
“(bb) the Chairman and ranking minority Member of the Committee on Environment and Public Works of the Senate; and
“(cc) the Chairman and ranking minority Member of the Committee on Natural Resources of the House of Representatives...

Figure 4: US H.R.6355 (115th)

While the bill elaborates of the “process”, the summary states that one was created. This type of summary would be hard to construct by a purely extractive method.

B.3 California Example

The California bills follow the same general patterns as US bills, but the format of some summaries is different. In Figure 5: the summary,

first, explains the existing law, then explains the change. The additional context is useful, and in the future we may build a system that references the existing law to create better summaries.

LEGISLATIVE COUNSEL'S DIGEST

Existing law requires that a bicycle operated during darkness upon a highway, a sidewalk where bicycle operation is not prohibited by the local jurisdiction, or a bikeway, as defined, be equipped with a red reflector on the rear that is visible from a distance of 500 feet to the rear when directly in front of lawful upper beams of headlamps on a motor vehicle. A violation of this requirement is an infraction.

This bill would instead require that a bicycle operated under those circumstances be equipped with a red reflector or a solid or flashing red light with a built-in reflector on the rear that is visible from a distance of 500 feet to the rear when directly in front of lawful upper beams of headlamps on a motor vehicle. By revising the definition of a crime, the bill would impose a state-mandated local program. The bill would also include a statement of legislative findings and declarations.

The California Constitution requires the state to reimburse local agencies and school districts for certain costs mandated by the state. Statutory provisions establish procedures for making that reimbursement.

Figure 5: California Bill Summary