
Topic Segmentation in the Wild: Towards Segmentation of Semi-structured & Unstructured Chats

Reshmi Ghosh^{1*} Harjeet Singh Kajal^{2†} Sharanya Kamath² Dhuri Shrivastava²

Samyadeep Basu^{3‡}

Soundararajan Srinivasan¹

1: Microsoft Corporation; 2: University of Massachusetts, Amherst; 3: University of Maryland, College Park

Abstract

Breaking down a document or a conversation into multiple contiguous segments based on its semantic structure is an important and challenging problem in NLP, which can assist many downstream tasks. However, current works on topic segmentation often focus on segmentation of structured texts. In this paper, we comprehensively analyze the generalization capabilities of state-of-the-art topic segmentation models on unstructured texts. We find that: (a) Current strategies of pre-training on a large corpus of structured text such as Wiki-727K *do not help* in transferability to unstructured texts. (b) Training from scratch with only a relatively small-sized dataset of the target unstructured domain improves the segmentation results by a significant margin.

1 Introduction

Topic Segmentation refers to the task of splitting texts into meaningful segments that correspond to a distinct topic or a subtopic. Natural language texts, especially in unstructured formats such as chat conversations and transcripts, often do not have an easy-to-detect separation between contiguous topics. Reliable & accurate division of text into coherent segments can help in making the text more readable as well as searchable. Hence, topic segmentation enables numerous applications such as search assistance and recommendation. It has also been noted that text segmentation can improve and speedup applications such as information extraction and summarization [7].

Recently, multiple approaches leveraging language models have been proposed for topic segmentation [7, 10, 4]. However, the datasets on which these approaches have been evaluated are often structured in nature such as Wiki-727K [7], Wiki-50 [7], RST, and Choi [3]. Adding to the constraints, in many applications, texts that need to be segmented are often unstructured such as chat transcripts and conversations. But, understanding the effectiveness of topic segmentation methods on such unstructured texts hasn't been well studied. In this paper, we empirically investigate the effectiveness of various topic segmentation methods on unstructured segmentation datasets such as LDC BOLT chat [14] and Amazon Topical chat [5]. In addition to being less structured than the Wiki-727K or Wiki-50 data, these datasets are challenging for traditional topic segmentation approaches as they contain grammatically ill-formed "noisy sentences," and a varying number of segments per conversation. Hence, we systematically examine the effectiveness of large-scale pre-training, dataset

*Corresponding Author; Email: reshminghosh@microsoft.com

†Work done in collaboration with Microsoft Corporation when a graduate student at UMass, Amherst

‡Work done while at Microsoft Corporation

used in pre-training, and fine-tuning strategies on these "out-of-domain" (data that is conversational in nature rather than the segmented Wiki content), unstructured text-segmentation datasets spanning traditional the LSTM-based models and the newer transformer-based architectures.

We find that large-scale pre-training (and fine-tuning with data from target domain) has only a *negligible* effect on downstream segmentation tasks, when the task consists of unstructured data. This is contrary to the conventional wisdom in NLP, where pre-training and fine-tuning is a common practice. We, therefore, identify topic segmentation on unstructured data as one task where large-scale pre-training doesn't have any significant effect. To perform well on segmentation of unstructured text, we find that training, from scratch, with only a *few* examples of the segmentation domain is sufficient. This is true for segmentation architectures ranging from the LSTM-based ones to the recent Transformer-based ones. Our results also show that, for unstructured topic segmentation, avoiding pre-training on a large corpus such as the Wiki-727K dataset results in saving a significant amount of training time and compute resources, facilitating the exploration of newer approaches. In summary, our contributions are as follows:

- We investigate and stress-test the effectiveness of current topic segmentation methods on unstructured texts, which is a more challenging segmentation task when compared to structured-segmentation tasks.
- We find that pre-training on large topic segmentation datasets such as Wiki-727K has negligible effect on downstream transfer to unstructured text-segmentation datasets and instantiating the model with only a few-examples of the unstructured task is sufficient.
- We provide simple and practical recipes to improve performance on unstructured-segmentation datasets.

2 Related Works

Topic segmentation has been explored through many realms; particularly, the approaches used could be broadly categorized as Non-neural-based and Neural-based approaches in both supervised and unsupervised settings.

Non-neural approaches: The early research efforts related to non-neural approaches for topic segmentation include [6] that focused on an unsupervised approach to analyze lexical cohesion in small segments by leveraging counts of word repetitions. This work was expanded to enable models understand words and sentences occurring in segments in a comprehensive manner leading to the wide use of lexical chains [1, 12].

Neural Approaches: [7] used a hierarchical Bi-LSTM to cast the topic segmentation as a supervised learning task. Other neural methods also leverage the Transformer architecture. In [10], the authors proposed three transformer-based models, of which Cross-Segment BERT model is particularly important for topic segmentation tasks as the model captures information from the local context surrounding a potential topic segment boundary to judge about which pool of sub-document units, the particular segment belongs. The other two model architectures use a hierarchical approach as used by [7], but with using the BERT model instead of BiLSTMs. [13] is another recent work that uses unsupervised approach based on BERT embeddings to segment topics in multi-person meeting transcripts.

3 Segmentation Datasets

Unstructured Datasets: In this work, We use the LDC's (Linguistic Data Consortium) BOLT SMS and chat data collection (filtered to only include chats and SMS in English; hereafter called as BOLT) and Amazon Topical Chat dataset. The BOLT dataset is significantly more unstructured in nature than Topical chat dataset, as the conversations in BOLT chat contain non-uniform sentence structure, incomplete sentences, and abbreviations. Thus, in this paper we refer to the BOLT data as unstructured dataset, and label the Topical chat dataset as semi-structured dataset (as it has more well-formed sentence structure; additional details in Appendix A.1).

Wiki-727K: Proposed by [7], this dataset contains 727,746 documents in the English language with text segmentations created according to the table of contents. The Wiki-727K dataset is considered a structured dataset because the text in it is not conversational in nature and has proper syntactical structure in the form of well-formed sentences grouped into paragraphs and sections.

These datasets were pre-processed (Appendix sub-section A.2) to generate labels for each sentence in every segment, for use in the supervised learning formulation.

Task	Datasets			Cross Segment BERT	Cross Segment RoBERTa-Base	Hierarchical Bi-LSTM
	Pre-Train	Finetune	Test			
A.1	Wiki-727K	-	Topical Chat	0.492	0.487	0.021
A.2	Wiki-727K	BOLT	Topical Chat	0.470	0.406	0.391
A.3	Wiki-727K	Topical Chat	Topical Chat	<u>0.725</u>	<u>0.713</u>	<u>0.931</u>
A.4	BOLT	-	Topical Chat	0.491	0.498	0.611
A.5	BOLT	Topical Chat	Topical Chat	0.734	0.729	0.915
A.6	Topical Chat	-	Topical Chat	0.764	0.767	0.951
A.7	Topical Chat	BOLT	Topical Chat	0.767	0.759	0.501
B.1	Wiki-727K	-	BOLT	0.487	0.467	0.005
B.2	Wiki-727K	BOLT	BOLT	<u>0.489</u>	<u>0.479</u>	<u>0.406</u>
B.3	Wiki-727K	Topical Chat	BOLT	0.511	<u>0.492</u>	0.152
B.4	BOLT	-	BOLT	0.569	0.561	0.443
B.5	BOLT	Topical Chat	BOLT	0.518	0.509	0.181
B.6	Topical Chat	-	BOLT	0.544	0.542	0.157
B.7	Topical Chat	BOLT	BOLT	0.536	0.529	0.331
C.1	Wiki-727K	-	Wiki-727K	0.604	0.599	0.57
C.2	Wiki-727K	BOLT	Wiki-272K	0.433	0.435	0.501
C.3	Wiki-727K	Topical Chat	Wiki-727K	0.513	0.509	0.411
C.4	BOLT	-	Wiki-727K	0.487	0.492	0.12
C.5	BOLT	Topical Chat	Wiki-727K	0.489	0.478	0.027
C.6	Topical Chat	-	Wiki-727K	0.505	0.502	0.020
C.7	Topical Chat	BOLT	Wiki-727K	0.5089	0.511	0.198

Table 1: Effect of model architectures and training strategies on topic segmentation tasks, grouped by the test dataset of choice - Topical Chat (A.1 - 1.7), BOLT (B.1 - B.7) & Wiki-727K (C.1 - C.7). The best F1 scores for **BOLT** (unstructured), **Topical Chat** (semi-structured), and **Wiki-727K**, per model are highlighted. Comparing the F1 scores in different pre-training & fine-tuning scenarios, we conclude that the Cross-Segment BERT model *consistently* outperforms the Hierarchical Bi-LSTM, and CSRoBERTa model on *almost* all tasks. We also find that large-scale pre-training with the structured Wiki-727K dataset (and then fine-tuning with data from Target domain - underlined; Task B.2 vs. B.4 and Task A.3 vs. A.6) is not required to create cohesive topic segments on unstructured or semi-structured conversational data.(Train-finetuning-test split described in Appendix A.3).

4 Experiments

We study the problem of segmenting semi-structured and unstructured chats using three popular modeling paradigms used in the structured segmentation domain: the Hierarchical Bi-LSTM model proposed by [7] and the Cross-Segment BERT model [10] (hereafter CSBERT), and the Cross-Segment RoBERTa (CSRoBERTa) [4] (Appendix A.4). In CSRoBERTa, we use the same training paradigm as CSBERT, but replace BERT with RoBERTa-base. Note that these models in the existing body of work have not been validated against any form of semi-structured or unstructured datasets. Table 1 presents the results of our analysis. We first pre-train all three models on Wiki-727K and test against the unstructured BOLT, the semi-structured Topical Chat, and the structured Wiki-727K datasets (underlined in Table 1). We then use different pre-training and fine-tuning combinations to examine the necessity of large-scale pre-training on structured datasets for segmenting conversational data. Also note that the results in Table 1 are generated by partitioning the BOLT and Topical Chat data into documents of 5 segments each (validated in Figure 1; more details on train-test split are in Appendix A.3).

Adapting the hierarchical Bi-LSTM model on the unstructured BOLT and the semi-structured Topical chat datasets during the validation phase, we conclude that the F1 scores (additional details on evaluation metrics in Appendix subsection A.5) are significantly worse when compared with the performance of the CSBERT and CSRoBERTa model in the same setting. However, evaluating the performance on the structured Wiki-727K dataset, we find that the F1 scores from both models are in the same range.

4.1 Effectiveness of Pre-training on Wiki-727K

In this section, we investigate the effectiveness of pre-training on the large Wiki-727K dataset. We first pre-train the Hierarchical Bi-LSTM, CSRoBERTa, and CSBERT models on the 80% of the Wiki-727K corpus. We further fine-tune the models on the unstructured datasets: BOLT chat and Topical Chat dataset. Additionally, we experiment with fine-tuning the models after training on semi-structured/unstructured dataset instead of the pre-training step with Wiki-727K. From Table 1, we find that training models with unstructured/semi-structured and fine-tuning it with semi-structured and unstructured dataset respectively, leads to better performance than fine-tuning with the Wiki-727K

checkpoint. For instance, we find that with the pre-train with Wiki-727K and fine-tune paradigm, using the CSBERT, we obtain a F1 score of 0.725 for Topical Chat, whereas only training on Topical Chat results in a F1 score of 0.764 (*Task A.3 vs. A.6*)⁴. For BOLT, the pre-train and fine-tune paradigm results in a F1 score of 0.489 whereas training from scratch with BOLT dataset results in a F1 score of 0.499 (*Task B.2 vs. B.4*). These results indicate that the conventional approach of pre-training on a large structured Wiki-727K dataset, and then fine-tuning with semi-structured or unstructured dataset, doesn't lead to an improvement in F1 scores, making the approach of pre-training on structured dataset questionable.

4.2 Effect of architecture on unstructured datasets

We test our initial conclusions (Table 1 on the efficacy of architectures across various number of segments. We find that, across different number of segments, the CSBERT model performs substantially better than the Hierarchical Bi-LSTM model and slightly better than CSRoBERTa model. From Figure 1, we conclude that:

- Across the varying number of segments curated and adapted during the pre-training step with Wiki-727K, the CSBERT model results in higher F1 scores when tested against semi-structured Topical Chat and unstructured BOLT datasets. Hence, we can conclude that for these topic segmentation tasks, CSBERT model is more suitable than the Hierarchical Bi-LSTM and the CSRoBERTa models.
- As the number of segments in the conversational datasets increases, the F1 scores from all models drop. We associate this finding with the fact that the large segments of the conversational data may contain more heterogeneous topics, making it difficult for all models to group coherent chats.

4.3 Practical recipes to improve unstructured segmentation tasks




Figure 1: Effect of model architecture and the number of segments on segmentation performance. Comparing the F1 scores from all three models pre-trained on Wiki-727K and inferred on unstructured & semi-structured datasets (without any fine-tuning), we find that CSBERT outperforms the other models robustly across the varying number of segments. Bi-LSTM performs very poorly compared to the other two models, which can be explained by its inability to fully consider the semantic context of the text representing segment boundaries. Additionally, with an increase in the number of segments, the F1 score peaks at 5 segments, and then drops across all models.

the rest with 0.2 yields the best results on these datasets. From Table 2, we conclude that re-weighting the cross-entropy loss function to provide proportional importance to both labels leads to a slightly better F1 score for all three models.

4.3.2 Focal loss as an alternative loss function for imbalanced topic segmentation

Focal loss has been used widely to mitigate the risks involved with class imbalance for tasks related to object detection [8], credit-card fraud detection [2], and other tasks involving class-imbalance [11]. We consider the α (a parameter that controls trade-off between precision and recall) and γ (focusing parameter; defines the degree of confidence assigned by the model to correct predictions that contributes to overall loss values) as focal-loss hyper-parameters and tune these over 10 epochs.

Casting topic segmentation for unstructured datasets as a binary classification problem leads to a severe imbalance in class labels, prompting the need to re-weight the samples or even modify the loss function to boost performance. The number of end-of-segment sentences (encoded as '1') is significantly smaller than the non-boundary sentences (sentences that do not mark the end of a segment; encoded as label '0') due to the inherent structure of segments in any document or chat.

4.3.1 Re-weighting in cross-entropy loss

To reduce the effect of dominance by the samples with label '0' and avoid biasing the model at inference time, we re-weight the class labels in cross-entropy loss function, giving proportional importance to labels '0' and '1'. The set of weights to be assigned is considered as a hyper-parameter and is optimized in the range [0, 1].

We find that weighting the end-of-segment sentences (encoded as '1') with 0.8, and weighting

⁴We find that for Cross-Segment BERT model Task A.6 and A.7 have very similar F1 scores

Datasets			Cross-Entropy Weights = [0.2, 0.8]			Focal Loss ($\alpha = 2; \gamma = 2$)		
Pre-Train	Finetune	Test	CSBERT	CSRoBERTa	Bi-LSTM	CSBERT	CSRoBERTa	Bi-LSTM
Wiki-727K	BOLT	BOLT	0.495	0.481	0.409	0.503	0.485	0.432
	-	BOLT	0.575	0.567	0.443	0.580	0.572	0.45
Wiki-727K	TC	TC	0.729	0.717	0.933	0.748	0.729	0.928
	-	TC	0.767	0.769	0.952	0.778	0.775	0.95

Table 2: Results from models trained with re-weighted cross-entropy loss and focal-Loss function. CSBERT refers to results obtained using the cross-segment BERT model. Additional results from all training strategies, as illustrated in Table 1 can be found in Appendix A.6, Table 3.

From Table 2, we find that the resultant F1 scores with focal loss are higher than re-scaling the cross-entropy loss function. Hence, we conclude that for future iterations of topic segmentation tasks involving unstructured datasets, focal loss is a better alternative.

5 Conclusion

In this work, we analyzed the effectiveness of current topic segmentation methods on unstructured texts. We find that, across different architectures and datasets, pre-training on the large, structured Wiki-727K *is not required* for segmentation of unstructured datasets such as BOLT and Topical Chat that contain syntactical and semantical noise. Training from scratch with only a few-examples provides a sufficiently strong baseline for this task. This finding is counter to the pre-training on a large corpus and fine-tuning on the target domain paradigm, commonly used in a variety of tasks in NLP. Further, we provided various practical fine-tuning recipes to boost the segmentation performance on such conversational and unstructured datasets.

References

- [1] Francisco Adarve, Helen Kwong, and Mike Speriosu. Topic segmentation of meetings using lexical chains. 2007.
- [2] Saptashwa Bhattacharyya. A loss function suitable for class imbalanced data: “focal loss” <https://towardsdatascience.com/a-loss-function-suitable-for-class-imbalanced-data-focal-loss-af1702d75d75>. 2021.
- [3] Freddy Y. Y. Choi. Advances in domain independent linear text segmentation. In *1st Meeting of the North American Chapter of the Association for Computational Linguistics*, 2000.
- [4] Goran Glavaš, Ananya Ganesh, and Swapna Somasundaran. Training and domain adaptation for supervised text segmentation. In *Proceedings of the 16th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 110–116, 2021.
- [5] Karthik Gopalakrishnan, Behnam Hedayatnia, Qinglang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, Dilek Hakkani-Tür, and Amazon Alexa AI. Topical-chat: Towards knowledge-grounded open-domain conversations. In *INTERSPEECH*, pages 1891–1895, 2019.
- [6] Marti A. Hearst. Texttiling: Segmenting text into multi-paragraph subtopic passages. *Comput. Linguist.*, 23(1):33–64, mar 1997.
- [7] Omri Koshorek, Adir Cohen, Noam Mor, Michael Rotman, and Jonathan Berant. Text segmentation as a supervised learning task. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 469–473, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.
- [8] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [9] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

- [10] Michal Lukasik, Boris Dadachev, Gonçalo Simões, and Kishore Papineni. Text segmentation by cross segment attention, 2020.
- [11] Jishnu Mukhoti, Viveka Kulharia, Amartya Sanyal, Stuart Golodetz, Philip Torr, and Puneet Dokania. Calibrating deep neural networks using focal loss. *Advances in Neural Information Processing Systems*, 33:15288–15299, 2020.
- [12] Laurianne Sitbon and Patrice Bellot. Topic segmentation using weighted lexical links (wll). In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 737–738, 2007.
- [13] Alessandro Solbiati, Kevin Heffernan, Georgios Damaskinos, Shivani Poddar, Shubham Modi, and Jacques Cali. Unsupervised topic segmentation of meetings with BERT embeddings. *CoRR*, abs/2106.12978, 2021.
- [14] Zhiyi Song, Dana Fore, Stephanie Strassel, Haejoong Lee, and Jonathan Wright. Bolt english sms/chat. <http://hdl.handle.net/10339/93641>. 2019.
- [15] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2017.

A Appendix

A.1 Dataset

To the best of our knowledge, Wiki-727K is the only publicly available large dataset, whose size makes it appropriate for large-scale pretraining. In fact, no singular conversational (unstructured) data of that scale exists, that can be used to pre-train large Transformer model. A similar approach of pre-training was undertaken in [7, 10, 4].

BOLT SMS and Chat dataset by LDC have conversations that have been extracted from messaging platforms like WhatsApp, iMessage, Android SMS, Symbian SMS, Viber, BlackBerry, QQ, Google chat, Skype chat, & Yahoo Messenger in languages like Chinese, Egyptian Arabic and English. The dataset contains 2140 Egyptian, 7844 Chinese and 9155 English conversations. We work with only the English portion of the data. Additionally, the dataset have the following characteristics:

- Mean sentence length (number of words) = 9.45 ; standard deviation = 9.41
- Mean segment length (number of sentences) = 11.28; standard deviation = 13.99

Topical Chat dataset from Amazon contains human to human conversations in 8 broad topics. It has more than 8000 conversations. Even though this dataset has well-formed sentence structure, it is in the form of conversations, and not articles, and thus we categorize the dataset as semi-structured. The underlying knowledge in Topical Chat spans across 8 broad topics and there is very little to no variation in the segment length of this dataset.

- Mean sentence length (number of words) = 19.86, standard deviation = 10.49
- Mean segment length (number of sentences) = 21.83, standard deviation = 1.75

A.2 Pre-Processing

Both Topical chat and BOLT datasets are conversational in nature, and thus required additional pre-processing. We achieve this by segmenting these datasets into multiple segments (we used 5 segments, as it resulted in best F1 scores; *Figure 1*), wherein each segment is a snippet of a specific chat conversation. Furthermore, we cast the problem of topic segmentation for OOD datasets as a supervised learning problem, and thus we create synthetic labels for each sentence in every segment. We first pre-process the datasets to synthetically create multiple segments, and label the sentences occurring in each segment $x_i = (x_1, x_2, \dots, x_n)$ according to the boundary condition, i.e., whether a specific sentence was end-of-segment (label, $y_i = '1'$) or non end-of-segment sentence (encoded label $y_i = '0'$).

A.3 Experimental Setup

For the initial task of pretraining on large structured dataset, we first randomly partitioned the Wiki-727K dataset in 80% / 10% / 10% format to create train, development, and test set. We then used train set from Wiki-727K dataset, to perform the pretraining using Hierarchical Bi-LSTM, Cross-Segment BERT, and Cross-Segment RoBERTa models.

Additionally for the initial task of pre-training on Wiki-727K and testing on semi-structured Topical Chat, and unstructured LDC dataset, we first partition all conversations to form segments (snippets of conversations) and group the segment chunks to form documents. The model at any point of the training process consumed a batch of these documents. As described in section A.2, the sentences in these segments are synthetically labelled to indicate the end-of-segment. As a result, the chunking the conversations from BOLT and Topical Chat datasets in 5 segments leads to 1815 and 1726 documents, respectively.

Furthermore, the documents created from the chunking process described above were split into train/fine-tuning/test sets for tasks described in Table 1 and Table 3. The documents from Topical Chat dataset was divided in 1099 / 348 / 281 documents for the train/fine-tuning/test splits respectively, wherein each document had 5 segments of mean segment length of approximately 11.5 sentences. Additionally, a similar approach was employed for the documents from BOLT dataset, and it was divided in 1090 / 363 / 362 documents for the train/fine-tuning/test splits respectively containing 5 segments each (mean segment length of 21.83 sentences).

Lastly to train the three models from scratch using the unstructured, and semi-structured data, without involving any pre-training on structured Wiki-727K, required 1109 documents of 5 segments from BOLT, and 1030 documents of 5 segments from Topical Chat.

A.4 Model Architecture

As detailed in Section 2, topic segmentation in the existing body of work has involved neural and non-neural approaches, but owing to the complexity of understanding heterogeneous conversational dataset, we use state-of-the-art neural models, i.e, Hierarchical Bi-LSTM as proposed by [7], and CSBERT model, introduced by [10].

The hierarchical Bi-LSTM model first learns the sentence representation and is then fed into a segment prediction sub-network. The lower-level sub-network has two-layer bi-directional LSTM layer that generates sentence representations by taking in words w_1, w_2, \dots, w_i of a sentence x_i as input to the model, and the intermediate output is passed through a max pooling layer to create the sentence representations e_i . The upper-level sub-network for segment prediction then takes a sequence of sentence embeddings generated from the previous layers, and feeds them to a two-layer Bi-LSTM layer, which then feeds into a fully connected-layer with softmax function to generate segmentation probabilities.

In the Cross Segment BERT model, the authors leverage information from the local context, i.e., studying the semantic shift in word distributions, as first introduced in [6]. The additional context on both sides of the segment boundary (termed as ‘candidate break in the paper), i.e., the sequence of word-piece tokens that come before and after the segment breaks. Basically the model is fed k word-piece tokens from the left and k tokens from the right of a segment break. The input is composed of a classification token (denoted by $[CLS]$), followed by the two contexts concatenated together, and separated by a separator token (denoted by $[SEP]$). The tokens are fed into the Transformer encoder([15]) which is initialized by $BERT_{LARGE}$ to output segmentation probabilities. The $BERT_{LARGE}$ model has 24 layers and uses 1024-dimensional embeddings and 16-attention heads. As the released BERT checkpoint only supports up to 512 tokens, we keep at maximum 250 word-tokens on each side.

The RoBERTa architecture was chosen as a comparative alternative to the BERT model in the Cross-Segment framework, as it is relatively newer successor of BERT [9]. However, rather than changing

the framework of Cross-Segment learning as proposed by [10], wherein the authors demonstrate the capability of BERT model to learn the context around end-of-segments in a robust way, we chose to utilize the same framework and simply replace the BERT model by RoBERTa-base.

A.5 Evaluation Metrics

Since we cast the task of topic segmentation as a binary classification problem, we use Precision, Recall and F1 scores to measure performance. Precision measures the percentage of boundaries identified by the model that are true boundaries. Recall measures the percentage of true boundaries identified by the model. Although comprehensive, it is important to note that there are some challenges associated with individually reporting Precision and Recall as they are somewhat less sensitive to near misses of true boundary identifications, when the prediction is off by one or two sentences, and thus we report F1 scores in Section 4. F1 score can be reliably used to conclude our initial findings from performing topic segmentation based binary classification on unstructured and semi-structured data.

A.6 Practical Recipes: Additional Results

In Table 3, we present our findings by experimenting with re-weighted cross-entropy loss and focal loss functions for all combinations of pre-training and fine-tuning strategies. We find that, across *almost* (except 3) all paradigms, Cross-Segment BERT model outperforms Cross-Segment RoBERTa and Hierarchical Bi-LSTM model on both re-weighting and focal loss function recipes. Hierarchical Bi-LSYTM F1 scores are largely in-consistent, and thus we can conclude the re-weighting of cross-entropy loss and replacement by focal loss function in our baseline model did not prove of large significance to the Bi-LSTM framework. We can attribute this finding to the RNN structure which limits contextual learning around the segment boundaries, whereas Transformer-based architectures, with their multi-head attention mechanism learn about the context robustly.

Moreover, Focal loss across all pre-training and fine-tuning strategies prove to be a robust alternative than re-weighting cross-entropy loss, for all three models.

Datasets			CE weights = [0.2, 0.8]			Focal Loss ($\alpha = 2; \gamma = 2$)		
Pre-Train	Finetune	Test	CSBERT	CSRoBERTa	Bi-LSTM	CSBERT	CSRoBERTa	Bi-LSTM
Wiki-727K	-	Topical Chat	0.497	0.491	0.028	0.512	0.504	0.037
Wiki-727K	BOLT	Topical Chat	0.475	0.411	0.398	0.483	0.427	0.405
Wiki-727K	Topical Chat	Topical Chat	0.729	0.717	0.933	0.748	0.729	0.928
BOLT	-	Topical Chat	0.493	0.5	0.613	0.501	0.510	0.606
BOLT	Topical Chat	Topical Chat	0.736	0.731	0.917	0.747	0.741	0.920
Topical Chat	-	Topical Chat	0.767	0.769	0.952	0.778	0.775	0.95
Topical Chat	BOLT	Topical Chat	0.768	0.761	0.505	0.777	0.768	0.515
Wiki-727K	-	BOLT	0.490	0.468	0.007	0.493	0.472	0.012
Wiki-727K	BOLT	BOLT	0.495	0.481	0.409	0.503	0.485	0.432
Wiki-727K	Topical Chat	BOLT	0.573	0.495	0.183	0.560	0.524	0.214
BOLT	-	BOLT	0.575	0.567	0.443	0.580	0.572	0.45
BOLT	Topical Chat	BOLT	0.520	0.511	0.183	0.531	0.519	0.189
Topical Chat	-	BOLT	0.546	0.544	0.159	0.555	0.551	0.164
Topical Chat	BOLT	BOLT	0.537	0.531	0.333	0.549	0.54	0.34
Wiki-727K	-	Wiki-727K	0.609	0.602	0.591	0.614	0.611	0.6
Wiki-727K	BOLT	Wiki-272K	0.435	0.438	0.508	0.441	0.446	0.513
Wiki-727K	Topical Chat	Wiki-727K	0.516	0.510	0.414	0.522	0.517	0.417
BOLT	-	Wiki-727K	0.489	0.495	0.127	0.493	0.501	0.132
BOLT	Topical Chat	Wiki-727K	0.492	0.479	0.03	0.496	0.484	0.037
Topical Chat	-	Wiki-727K	0.507	0.503	0.023	0.516	0.510	0.031
Topical Chat	BOLT	Wiki-727K	0.511	0.514	0.199	0.520	0.521	0.210

Table 3: Comparing the F1 scores with different pre-training & fine-tuning scenarios using re-weighting and Focal Loss, we conclude that Focal loss stands out to be a much a robust alternative than re-weighted cross-entropy loss for topic segmentation of unstructured texts. The best F1 scores for **BOLT**, **Topical Chat**, and **Wiki-727K**, per model are highlighted.