

# A New Dataset for Topic-Based Paragraph Classification in Genocide-Related Court Transcripts

Miriam Schirmer<sup>†,‡</sup>, Udo Kruschwitz<sup>†</sup>, Gregor Donabauer<sup>†</sup>

<sup>†</sup>University of Regensburg, Germany

<sup>‡</sup>TUM School of Social Science and Technology, Technical University of Munich, Germany  
{miriam.schirmer, udo.kruschwitz, gregor.donabauer}@ur.de

## Abstract

Recent progress in natural language processing has been impressive in many different areas with transformer-based approaches setting new benchmarks for a wide range of applications. This development has also lowered the barriers for people outside the NLP community to tap into the tools and resources applied to a variety of domain-specific applications. The bottleneck however still remains the lack of annotated gold-standard collections as soon as one’s research or professional interest falls outside the scope of what is readily available. One such area is genocide-related research (also including the work of experts who have a professional interest in accessing, exploring and searching large-scale document collections on the topic, such as lawyers). We present GTC (Genocide Transcript Corpus), the first annotated corpus of genocide-related court transcripts which serves three purposes: (1) to provide a first reference corpus for the community, (2) to establish benchmark performances (using state-of-the-art transformer-based approaches) for the new classification task of paragraph identification of violence-related witness statements, (3) to explore first steps towards transfer learning within the domain. We consider our contribution to be addressing in particular this year’s hot topic on Language Technology for All.

**Keywords:** Text classification, BERT, Professional Search, Genocide Studies, Language Technology for All

## 1. Introduction

Information overload has led to a multitude of search applications of which Web search is just one out of many. Unlike search for leisure or personal interest there is a vast area of search contexts which are found in a work environment. Professional search falls into that scope, i.e. search over domain-specific document collections and often with search tasks that are recall-oriented rather than precision-focused (Kruschwitz and Hull, 2017; Verberne et al., 2019). Beyond applications where such search effort can directly be measured in financial terms (e.g. in patent search, e-discovery or the compilation of systematic reviews) there are many other fields where these costs are more implicit, e.g. in the area of genocide studies that rely on the analysis of vast quantities of different resources (Bachman, 2020; Hinton, 2012).

Looking at the wider picture, searching large text corpora for specific thematic patterns can be very time-consuming and non-trivial, in particular for searchers who do not have a solid foundation in NLP or search technology. The huge amount of court transcripts of genocide tribunals presents a perfect example: the *International Criminal Tribunal of the Former Yugoslavia (ICTY)* alone provides official transcripts for each of its cases online, leading up to approximately 2.5 million pages of transcripts in total (ICTY, 2016). Searching for specific content in a text corpus like this usually requires vast amounts of manual research capacity (Hoang and Schneider, 2018). Tools and approaches to augment this type of search and help limit manual efforts have been developed for a broad

range of use cases, e.g. for automating search strategies or text extraction from documents (MacFarlane et al., 2021; Russell-Rose et al., 2021). However, even with the help of suitable tools, searching for specific text passages in large text corpora generally remains a difficult task, in particular when the search is recall-oriented (Bache, 2011; Kaptein et al., 2013; Noor and Bashir, 2015). It should also be pointed out that in many use cases in which experts have to sift through large amounts of textual data a *fully automated* analysis might neither be achievable nor desirable and the provision of support tools that *assist* the expert are the preferred option. The area of fact-checking is one such application context (Nakov et al., 2021).

Turning to our own use case, the search for specific content in transcripts of genocide tribunals further proves difficult because transcripts are only accessible individually (usually one court day per transcript) and in different formats, depending on the tribunal. So far, no datasets of any kind containing genocide court transcripts have been published. Similarly, no other forms of pre-structured or annotated text data in this field of research exist.

This paper addresses this gap by providing a systematically annotated dataset containing text material from three different genocide tribunals: the *Extraordinary Chambers in the Courts of Cambodia (ECCC)*, the *International Criminal Tribunal for Rwanda (ICTR)*, and the *ICTY*. In addition to compiling the sampled corpus, we provide annotations within the text. More specifically, text passages in which witnesses talk about experienced violence have been annotated, focusing on a core part of each testimony. Given that respective

passages on violence often cover crimes that are relevant for the indictment of the accused, such as murder or rape, they are essential for judgement. At the same time, they are not easily classifiable due to their potential ambiguity.

Since this dataset of textual documents is the first in this area, we hope that it provides a valuable resource for NLP-based genocide research. To foster generalisability and assess transferability of approaches, the corpus contains a sample from different tribunals.

We consider the provision of the new resource our main contribution, but we also provide experimental work that will serve as a benchmark and allow the contextualisation within the broader field. We use fine-tuned BERT models for that purpose. Utilizing the heterogeneous nature of the corpus we also explore transfer learning and report results.

Beyond the contribution to the NLP community, it is our hope that the results of this paper will be useful for both scholars and practitioners at international criminal tribunals who need to work through large quantities of transcript material as part of their everyday job.

We summarize our contributions as follows:

1. We present GTC, a new reference corpus sampled from different international criminal courts in the context of genocide tribunals. The corpus contains annotations of statements by witnesses about experienced violence.
2. We built state-of-the-art transformer-based classifiers to provide benchmarks for the new classification task of paragraph identification of violence-related witness statements.
3. We provide experimental results for transfer-learning by varying the training and testing data across documents from different tribunals.
4. We make all data as well as code available to the community.<sup>1</sup>

## 2. Related Work

We touch on the three key areas of interest our work falls into, namely resources, professional search, and text classification. The discussion of each of these should simply serve as both a motivation and basic context.

### 2.1. Resources

The importance of publicly available language resources to help develop NLP applications has long been recognized, e.g. Calzolari et al. (2010), and the domain-specific nature of many problems is what makes repositories such as the *LRE Map*<sup>2</sup> a valuable starting point for many researchers and practitioners.

<sup>1</sup><https://github.com/MiriamSchirmer/genocide-transcript-corpus>.

<sup>2</sup><https://lremap.elra.info>

For the specific use case of assisting searchers to identify relevant information in genocide-related court transcripts resources are very limited to non-existent. Of course, transcripts of each tribunal are available through the respective courts' websites – however, their quality in terms of digitisation (e.g., object character recognition) varies greatly. Specifically for the ICTY, Fidahić (2021) further criticises that transcripts are only available in certain languages, thus limiting access mainly to English-speaking readers. Considering that the field of genocide research and studies (Totten, 2017) is multi-faceted enough to warrant the provision of suitable resources, we see our own contribution as a starting point to fill this gap, even though we limit our work to English transcriptions.

### 2.2. Professional Search

Searching through court transcripts can often be framed as an instance of professional search (Koster et al., 2009; Russell-Rose et al., 2021). Professional search describes the process of searching for information in a work context which is commonly domain-specific and requires expertise in a specific area. Key features of professional search are limited time and budget resources, making it desirable to provide support that helps classify specific text passages which ultimately could drastically reduce search efforts (Russell-Rose et al., 2021). It should be noted that professional search is very different from other types of search such as Web search. A common observation is that searches take a lot longer to satisfy a specific information need. For example, Bullers et al. (2018) found that librarians spend 26.9 hours on average on systematic reviews that involve searching for specific content, indicating that this task is highly time-consuming. Similarly, Greene-Colozzi et al. (2021) discuss the time-consuming process of researching court transcripts and other relevant sources related to cybercrime. Professional search in court transcripts in general, however, has not been analysed so far.

Another important aspect dealing with extensive search in large text corpora are human factors. Especially when dealing with time-consuming search in text documents that lasts for hours, fatigue might be an issue that reduces the quality of the search. Additionally, manual search is also more vulnerable to subjectivity, motivating the use of automated search algorithms (Li et al., 2020).

In the context of professional and augmented search, different supported search scenarios could also be helpful as a first step. Especially when working with an annotated dataset that is built around a binary classification task, the classification labels provided can help to significantly narrow down the text material for further in-depth search.

Different tools and algorithms to save time in searching through text have been discussed – varying strongly depending on the specific search context. For example,

MacFarlane et al. (2021) give a broad overview of different tools for systematic literature reviews, such as tools for text and data extraction or automatic query expansion. While these tools might help review literature more efficiently, the authors note that their use is not widespread. Furthermore, not all of the above-mentioned tools are helpful when it comes to content-based search in text documents. In this context – when looking for specific content in court transcripts – tools for enhanced keyword search might prove more useful.

### 2.3. Text Classification

Topic-based paragraph classification specifically for court transcripts has so far not been discussed in the literature. Nevertheless, extensive research in this area has been done in other fields. As a traditional and fundamental NLP task, text classification covers a wide range of tasks ranging from category labeling over sentiment analysis to authorship attribution (Jurafsky and Martin, 2021). Traditionally effective approaches to supervised machine learning, such as Support Vector Machines (Al Amrani et al., 2018; Tong and Koller, 2001; Zhang et al., 2007) or K-nearest neighbour (Bijalwan et al., 2014), have now largely been replaced by transformer-based approaches (Dhar et al., 2021; Jurafsky and Martin, 2021; Minaee et al., 2021). Fine-tuning BERT has become the standard baseline in text classification (Devlin et al., 2019), not just beating traditional machine learning paradigms but also recurrent neural networks (RNNs), convolutional neural networks (CNNs) or other deep neural networks (DNNs) (Li et al., 2020).

Example topic areas in which BERT has been utilized effectively in text classification include various forms of sentiment analysis ranging from aspect-based sentiment analysis (Sun et al., 2019) to sentiment analysis on the impact of coronavirus in social life (Singh et al., 2021), as well as reading comprehension tasks, e.g., Xu et al. (2019).

Of specific concern to our underlying use case is text classification that requires *text segment classification*, commonly found when applied to social media data, such as tweets and comments. For this type of analysis, splitting larger text data into paragraphs limited to a certain number of words has been established as a regular step in the NLP pipeline (Li et al., 2020). A very prominent example of using BERT sequence classification is hate speech detection (e.g., Mozafari et al. (2020a), Mozafari et al. (2020b), Sohn and Lee (2019)). By applying BERT to Twitter data, tweets can easily be classified according to whether they contained racism, sexism, or hate, among others (Mozafari et al., 2020a).

New BERT models and applications are being reported at rapid speed as the model is continuously applied in new fields. Examples are the recently developed ClimateBert, a pre-trained language model for climate-related text (Webersinke et al., 2021) or COVID-

Twitter-BERT (CT-BERT) (Müller et al., 2020).

### 2.4. Concluding Remarks

The new corpus we provide aims to bridge a (domain-specific) gap that exists in the landscape of annotated text collections. In order to assess the utility of the corpus and the difficulty of the underlying classification task we will adopt the commonly applied baseline approach of fine-tuning BERT. One of the goals is to show whether or not BERT also serves as an efficient tool for this type of text data and whether it can help simplify classification of paragraphs in court data.

This can only be a first step at filling the identified gap – there will be scope for many future directions, not least to replicate the approach to other languages.

## 3. Genocide Transcript Corpus (GTC)

We introduce *Genocide Transcript Corpus (GTC)*, a corpus of transcripts drawn from the court proceedings of international tribunals dealing with cases of genocide. Following sampling of the original data we also apply an annotation step that assigns binary labels to individual paragraphs. The paragraph labeling is aimed at identifying those parts of the text that refer to violence experienced by witnesses – relevant are only those text segments which are actually part of witness statements.

The dataset used in this study consists of 1475 text passages from three different genocide tribunals. Transcripts from the three biggest ad-hoc genocide tribunals, the ECCC, the ICTR, and the ICTY were selected. In a first step, the courts' databases were searched for witnesses who have actually experienced some form of violence. This pre-selection ensured having a substantial amount of relevant text passages in the dataset and thereby excluding technical or expert witnesses. Three different tribunals were selected to provide a diverse dataset and explore transfer learning, i.e. to show possible differences in the results after training and testing with data from different tribunals. Thus, results are more generalisable and differences in individual tribunals are controlled for.

Between 4 and 7 transcripts were selected per tribunal and were divided into equally large text chunks of 250 words each. Numbers and punctuation were removed in a first preprocessing step. In the final dataset, the number of samples is roughly equally distributed across tribunals (ECCC: 465, ICTY: 530, ICTR: 480). Differences occur since only complete transcripts with varying length (about 40 to 120 pages) were included.

The current version of the GTC contains the following data:

- For the ECCC, transcripts with a total of 438 pages from two different trials (Case 001 against Kaing Guev Eav, Case 002 against Nuon Chea and Khieu Samphan) were selected. This includes the

proceedings of 4 full court days and the hearing of 7 witnesses.

- Transcripts of the ICTY were taken from the cases against Slobodan Milošević (IT-02-54) and Duško Tadić (IT-94-1). The material consists of 416 pages of transcripts from 5 trial days, with 15 witnesses testifying in court.
- For the ICTR, 566 pages of transcript material from the cases against Jean-Paul Akayesu (ICTR-96-04) and Pauline Nyiramasuhuko et al. (ICTR-98-42) were included in the dataset. The ICTR data includes 5 witnesses and 7 court days.

In total, 1420 pages of transcripts were incorporated into the dataset. Differences in the number of pages and witnesses are firstly due to different transcript formats regarding digitisation and text density per page. Secondly, legal proceedings vary between the different tribunals and thus lead to slightly different content. For example, in the selected ECCC and ICTR transcripts, witnesses are questioned for approximately one court day, whereas in the selected ICTY transcripts, 2 to 3 witnesses were questioned per day.

## 4. Methodology

### 4.1. Label Annotation

All samples were labeled according to whether they contain a witness's description of experienced violence (0 = no violence, 1 = violence). Violence in this context is interpreted broadly and includes accounts of experienced or directly witnessed torture, interrogation, death, beating, psychological violence, experienced military attacks, destruction of villages, looting, and forced displacement. We restrict our interest to a binary classification, i.e. different acts of violence were not categorized further into subcategories. Figure 1 provides an example of a rather clear distinction between the two labels.

An important requirement for labeling text passages as containing accounts of violence was whether experienced violence was described by the witness orally in court. Questions by lawyers and judges containing violence-related words were thus labeled '0'. However, since the words used in both cases are the same for the most part, the differentiation between violence-related statements of witnesses vs. lawyers, judges, or the accused makes an automated classification more difficult. Having written statements (e.g., statements recorded previously by court staff, police, or human rights organisations) read out loud during the trial increases this difficulty further: even though reports contain accounts of experienced violence, they are not labeled '1' because they were not expressed orally by the witnesses during the trial, but by a lawyer or another representative of the court (see Figure 2 for an example).

It should have become clear that the task of correct classification in an automated fashion is non-trivial;

#### Label 0

Q. [...] As we discussed before, I will ask you some questions concerning your experiences in Rwanda back in 1994. Back in April of 1994 where did you live? And please you can just specify by commune.

A. We were living in Taba commune.

Q. Is that in Rwanda?

A. It's a commune in Rwanda, in Gitarama prefecture.

Q. Around the beginning of April did you ever receive news of the crash of the president's plane?

A. Yes, I heard this. [...]

ICTR-96-4-I, October 23<sup>rd</sup> 1997, p. 17-18.

#### Label 1

Q. What happened next?

A. He took me and he had a very long knife that he was wearing in his belt and also a small ax in his hand. We arrived near the primary school. The classrooms are very close to the bureau communal, very close to the place where we were before and it's very close to the road, as well, and when we arrived at that location **this child put down this ax, he also put down the long knife, near me, and you see these things are not very easy to see, a young child like that rape me.** I hope you understand that this is something that is very, very painful. [...]

ICTR-96-4-I, October 23<sup>rd</sup> 1997, p. 60.

Figure 1: Sample abstracts from the corpus demonstrating two clear-cut examples for a text passage that does *not* contain accounts of violence in a witness statement (top example – Label 0) and one that does (bottom example – Label 1). The examples were shortened, and both format and punctuation were adapted for readability.

simple 'bag of words'-based approaches are likely to underperform. Apart from the context that makes it clear how to classify a paragraph, looking at the vocabulary alone will not be sufficient. A similar observation was made when classifying a corpus of tweets which were classed as falling into a number of different classes all to do with violence such as crises, violence, accidents, and crime (Alhelbawy et al., 2016). It was found that the inter-rater agreement varied significantly across the different violence classes.

Since this dataset does not differentiate between subcategories, classification was limited to a binary task. However, to make sure that the categorization is reliable, a random selection of approximately 200 samples were independently labeled by a second researcher (with an inter-rater reliability  $\kappa = 0.86$ ) according to the above-mentioned facets of experienced violence. Even though only a sample of the dataset is labelled by two annotators, the high inter-rater agreement suggests that

**Label 0 (introduction of witness by lawyer)**

[...] The witness is a journalist working for a newspaper and he has reported several materials during the conflict in 1998. [...] He describes the situation in Suva Reka on the 25th of March, 1999, including the **killings and burning of houses**. [...] The witness also describes that on April 1st, Belanica was shelled, and police, military, and paramilitary forces, numbering about 1.500, subsequently entered the village. **The Serb forces forced people from their houses, looted their homes, loaded the goods on the trucks, and set the houses on fire.**

ICTY, 020424IT, April 24<sup>th</sup> 2002, p.3361-3362.

Figure 2: Example of a text passage that contains violence-related vocabulary, but is not labeled 1. As in Fig. 1, this example was shortened and adapted.

the labeling process yielded sufficiently plausible results.

Table 1 provides an overview of the number of each label per tribunal. Differences in the label balance are due to the random selection of transcripts.

	$n_0$	$n_1$	$n_{total}$
ALL	946	529	1475
ECCC	286	179	465
ICTY	401	129	530
ICTR	259	221	480

Table 1: Overview of label balance for the complete dataset (“ALL”) and the three individual tribunal datasets.

## 4.2. Experimental Setup

For all experiments, the 12-layer BERT<sub>base</sub> architecture for sequence classification (Devlin et al., 2019) was used to classify text passages of genocide tribunal transcripts.

As described in Section 3, the dataset consists of 3 subsets with data from different tribunals. 5-fold cross-validation (80:10:10) was applied to each subset and to the full version of the dataset (concatenated subsets).

Overall, BERT<sub>base</sub> was trained on all possible train, validate and test constellations, leading to a total of 16 different combinations. In those cases, in which training, validation and test data originate from the same subset, the respective splits led exactly to an 80:10:10 distributed number of samples. When training on one (or more) class(es) and testing on samples of a single remaining class, we held out all samples of the target class for testing. Consequently, for some of the combinations the number of test samples equals or even exceeds the number of samples in the train and validation data (for details see Table 2).

In a first step, BERT<sub>base</sub> was trained on the full dataset to classify samples of all three tribunals, but also to classify tribunal-specific text chunks.

Secondly, we apply the same setup to all three subsets. More specifically, training was performed using tribunal-specific samples to see if BERT is still able to predict class labels of both, the mixed dataset (excluding training class), as well as the remaining tribunal-specific subsets.

To test for the detection of undersampled violence-related paragraphs, additional experiments on this data were set up. All of the subset-specific negative class samples were used and a random proportion of 20% of positive class samples was added.

For training and validation a batch-size of 16 samples and an epoch-number of 3 (compare Devlin et al. (2019)) was used. The training was executed using 4 Nvidia RTX 2080Ti GPUs with an overall memory size of 44GB.

Precision, recall, micro and macro F1 scores for each train/validate/test constellation are provided – in line with common practice, macro F1 scores will be the reference score when comparing results (Jurafsky and Martin, 2021).

## 5. Results

Our results show that a binary classification based on BERT yields very reliable results across text data from different tribunals. A macro F1 score of 0.81 when training, testing and validating with the complete, mixed dataset that includes all three tribunals shows that BERT can be applied to this type of data and provides reasonably good predictions across the different subsets.

Considering the individual tribunals, using a tribunal-specific dataset for training and validating provided varying test results (ECCC-ECCC macro F1=0.70; ICTY-ICTY macro F1=0.68; ICTR-ICTR macro F1=0.80). Overall, using the mixed dataset for training and validating resulted in the highest F1 scores throughout the tribunal variations (min macro F1=0.78, max macro F1=0.85), independently of the dataset that was used for testing. The highest individual F1 score in our experiments was obtained when predicting data from ICTR transcripts with trained and validated data from the mixed dataset (“ALL”) (macro F1=0.85).

Looking at the tribunal-specific outcomes for the respective training/validating/sets also yielded solid results overall: Interestingly, using the ECCC data for training and validating has the highest true prediction rates when testing is conducted with ICTR data (macro F1=0.79), whereas using ECCC data for training, validating *and* testing only led to a comparatively low macro F1 score of 0.70. When training with ICTY data, performance was also best when predicting ICTR data (macro F1=0.81). Results are similar for training and validating with ICTR data: The highest macro F1 score (0.80) was obtained when using ICTR data for testing.

Train/val data	Test data			
Mixed dataset	Mixed dataset ( $n_{train}=1180$ , $n_{val}=147$ , $n_{test}=148$ )	ECCC dataset ( $n_{train}=808$ , $n_{val}=202$ , $n_{test}=465$ )	ICTY dataset ( $n_{train}=756$ , $n_{val}=189$ , $n_{test}=530$ )	ICTR dataset ( $n_{train}=796$ , $n_{val}=199$ , $n_{test}=480$ )
ECCC dataset	Mixed dataset ( $n_{train}=372$ , $n_{val}=93$ , $n_{test}=1010$ )	ECCC dataset ( $n_{train}=372$ , $n_{val}=46$ , $n_{test}=47$ )	ICTY dataset ( $n_{train}=372$ , $n_{val}=93$ , $n_{test}=530$ )	ICTR dataset ( $n_{train}=372$ , $n_{val}=93$ , $n_{test}=480$ )
ICTY dataset	Mixed dataset ( $n_{train}=424$ , $n_{val}=106$ , $n_{test}=945$ )	ECCC dataset ( $n_{train}=424$ , $n_{val}=106$ , $n_{test}=465$ )	ICTY dataset ( $n_{train}=424$ , $n_{val}=53$ , $n_{test}=53$ )	ICTR dataset ( $n_{train}=424$ , $n_{val}=106$ , $n_{test}=480$ )
ICTR dataset	Mixed dataset ( $n_{train}=484$ , $n_{val}=96$ , $n_{test}=995$ )	ECCC dataset ( $n_{train}=384$ , $n_{val}=96$ , $n_{test}=465$ )	ICTY dataset ( $n_{train}=484$ , $n_{val}=96$ , $n_{test}=530$ )	ICTR dataset ( $n_{train}=484$ , $n_{val}=48$ , $n_{test}=48$ )

Table 2: Overview of sample balance for the complete, mixed dataset (“ALL”) and the three individual tribunal datasets for each train/validate/test constellation.

Overall, precision and recall turned out to be fairly balanced throughout the different training and testing processes. See Table 3 for a detailed overview of the results. When conducting the experiments with under-sampled violence-related data, results turn out to be different. Despite using class weights for training (due to the underrepresented positive label), the results obtained are much lower than those reported for the full dataset. For each subset (ECCC: macro F1=0.51, micro F1=0.81; ICTY: macro F1=0.45, micro F1=0.74; ICTR: macro F1=0.45, micro F1=0.75) as well as for the mixed dataset (macro F1=0.47, micro F1=0.77) macro F1 scores are about half of the values reported so far. Since positive samples are heavily underrepresented (e.g. 1 out of 31 samples in the test set) precision, recall and binary F1 for this class amount to 0.0 for a range of data splits. This leads to the overall poor results for this setup. It also offers directions for future experiments.

## 6. Discussion

**General Discussion:** This study presented a new type of dataset for NLP-based research in the field of genocide and violence studies.  $BERT_{base}$  was further used to predict if text passages from court transcripts of three different genocide tribunals contain accounts of experienced violence by the respective witnesses. The results, in line with expectations, indicate that the mixed dataset is most successful when predicting if a certain text passage from one of three genocide tribunals contains accounts of experienced violence by a witness. Even when classifying paragraphs of one specific tribunal (e.g., the ECCC) with the model that was trained with data from the same tribunal (ECCC in this case), the model trained on the complete dataset provides better results. Including additional data from other tribunals thus improves the quality of the classification.

cation.

**Contextualisation:** Looking at the wider picture, binary classification scores vary widely across NLP applications (Arase and Tsujii, 2019; Wang et al., 2019) – direct comparisons with other studies must therefore be interpreted with caution. Nonetheless, the ballpark figures we obtained are comparable to state-of-the-art (BERT-based) performance on some other commonly used binary classifications such as MRPC (Zhang et al., 2021), but fall short of performance levels expected for other settings (d’Sa et al., 2020). On the one hand, this confirms once more that BERT can be successfully applied to our corpus and perfectly presents how well this language model has been developed in recent years. On the other hand, further fine-tuning will be necessary to solve performance-related shortcomings.

**Precision vs. Recall:** The overall similarity of precision and recall rates in our dataset implies that this type of classification might be useful for a broad range of applications. In some cases, recall rates might be more important than precision rates: for example, similar to patent search (Bache, 2011; Bashir and Rauber, 2010), a high recall is especially important when avoiding missed positive classifications is crucial. In a genocide-transcript-related context, this could apply to staff members who have to work through court transcripts as part of their daily work routine, e.g. for preparing a case. For this option specifically, applying the classification algorithm reduces the time spent on manual search drastically, making sure that no sample is missing and leaving time for manual adaptations. On the other hand, in the context of fast and efficient search with less time for manual adaptations, high precision rates would be more useful, e.g., when only some examples of relevant text segments are required and correctness is more important than completeness (Kong and Allan, 2016).

	ALL				ECCC				ICTY				ICTR			
	P	R	mac. F1	mic. F1	P	R	mac. F1	mic. F1	P	R	mac. F1	mic. F1	P	R	mac. F1	mic. F1
<b>ALL</b>	0.81	0.83	0.81	0.83	0.81	0.82	0.82	0.82	0.78	0.78	0.78	0.83	0.85	0.85	0.85	0.85
<b>ECCC</b>	0.77	0.77	0.77	0.79	0.77	0.72	0.70	0.75	0.73	0.71	0.71	0.78	0.81	0.79	0.79	0.80
<b>ICTY</b>	0.77	0.78	0.77	0.78	0.77	0.78	0.77	0.78	0.70	0.73	0.68	0.74	0.81	0.81	0.81	0.81
<b>ICTR</b>	0.74	0.74	0.74	0.78	0.79	0.77	0.78	0.79	0.69	0.74	0.70	0.75	0.83	0.78	0.80	0.85

Table 3: Results for macro precision (P), macro recall (R) and macro/micro F1 scores on test data (columns) with respect to different training/evaluation set (rows) combinations.

### Number of Text Chunks and Labeling Balance:

When looking at the results of this study, imbalances in the number of text chunks, labels and in the train-validate-test-ratio must be kept in mind. Still, in spite of the ICTY data containing fewer violence-related text segments, results did only differ slightly, indicating that this label imbalance does not impact the results significantly.

However, extending the dataset further would be a first step in making the results more stable. More text data could also help to improve the label balance: by selecting more transcripts per tribunal, the chances of choosing transcripts that contain no/few or above-average accounts of violence can be reduced.

An extended dataset would also make it easier to experiment with undersampled violence-related paragraphs. As already mentioned in Section 5, this setup currently lacks a sufficient number of positive labels in the test sets (when undersampling this class in an adequate ratio to keep the overall number of samples stable). Thus, adding more (non-violent) text chunks would make it easier to generate representative training/validation/test splits with a sufficient number of paragraphs for both classes regarding this setup. However, the dataset as it is offers directions for a range of possible experiments including the identification of violence related text chunks when heavily underrepresented.

**Future Research:** This dataset has the potential of serving as a basis for a variety of research approaches in the field of genocide research in the future. For example, more in-depth comparisons between linguistic or content-based characteristics between the three tribunals could be made, building bridges between the interdisciplinary field of genocide research and NLP-approaches. Since the provided dataset is violence-based, further research could, for example, build on psycho-linguistic aspects of violence-related trauma in witness statements of genocide tribunals.

From an NLP perspective, next steps could include further fine-tuning of BERT and establishing a model version that is pre-trained specifically on court transcripts of genocide tribunals. Conducting the experiments with more recent transformer architectures or machine learning techniques could also yield interesting results and would therefore be a good starting point for future research. Given that the full annotated dataset is

publicly available online, further studies could also include a detailed error analysis of the misclassified paragraphs. Not least we see our work as a first step towards a downstream practical search system.

## 7. Conclusion

This paper introduces a new dataset of genocide transcript data as a basis for further NLP research and applications. In addition, a baseline for classifying the transcript samples into violent or non-violent text chunks respectively is provided. The results, based on the well-established BERT architecture, demonstrate that such models can successfully be applied to this new domain and its related classification task. Although the number of text segments used in this study could be further extended (as it especially was observable during experiments with undersampled violence related paragraphs), classification with BERT proved to be successful, emphasizing once more the potential this language model holds even for research areas that have not been in the focus of NLP applications.

## 8. Ethical Considerations

All of the transcripts used in this paper are published online on the respective courts' websites and are publicly accessible. Since this type of text material contains personal and highly sensitive information about witnesses before international criminal courts, special care was taken to ensure that text fragments were not taken out of context. The use of witness names (or their anonymisation) in the dataset was adopted according to the original court document.

## 9. Acknowledgements

This project was supported by the Foundation of German Business (Stiftung der deutschen Wirtschaft e.V., sdw). We would further like to thank the three anonymous reviewers for their helpful comments.

## 10. Bibliographical References

Al Amrani, Y., Lazaar, M., and El Kadiri, K. E. (2018). Random forest and support vector machine based hybrid approach to sentiment analysis. *Procedia Computer Science*, 127:511–520.

- Alhelbawy, A., Poesio, M., and Kruschwitz, U. (2016). Towards a corpus of violence acts in arabic social media. In Nicoletta Calzolari (Conference Chair), et al., editors, *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1627–1631, Portorož, Slovenia, May. European Language Resource Association (ELRA).
- Arase, Y. and Tsujii, J. (2019). Transfer fine-tuning: A BERT case study. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5393–5404, Hong Kong, China, November. Association for Computational Linguistics.
- Bache, R. (2011). Measuring and improving access to the corpus. In Mihai Lupu, et al., editors, *Current Challenges in Patent Information Retrieval*, pages 147–165. Springer, Berlin, Heidelberg.
- Bachman, J. (2020). Cases studied in genocide studies and prevention and journal of genocide research and implications for the field of genocide studies. *Genocide Studies and Prevention: An International Journal*, 14(1):2–20.
- Bashir, S. and Rauber, A. (2010). Improving retrievability of patents in prior-art search. In Cathal Gurrin, et al., editors, *ECIR 2010: Advances in Information Retrieval*, pages 457–470, Milton Keynes, UK, March. Springer.
- Bijalwan, V., Kumar, V., Kumari, P., and Pascual, J. (2014). KNN based machine learning approach for text and document mining. *International Journal of Database Theory and Application*, 7(1):61–70.
- Bullers, K., Howard, A. M., Hanson, A., Kearns, W. D., Orriola, J. J., Polo, R. L., and Sakmar, K. A. (2018). It takes longer than you think: Librarian time spent on systematic review tasks. *Journal of the Medical Library Association: JMLA*, 106(2):198–207.
- Calzolari, N., Soria, C., Gratta, R. D., Goggi, S., Quochi, V., Russo, I., Choukri, K., Mariani, J., and Piperidis, S. (2010). The LREC map of language resources and technologies. In Nicoletta Calzolari (Conference Chair), et al., editors, *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta, May. European Language Resources Association (ELRA).
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Dhar, A., Mukherjee, H., Dash, N. S., and Roy, K. (2021). Text categorization: Past and present. *Artificial Intelligence Review*, 54(4):3007–3054.
- d'Sa, A. G., Illina, I., and Fohr, D. (2020). BERT and fastText embeddings for automatic detection of toxic speech. In *2020 International Multi-Conference on "Organization of Knowledge and Advanced Technologies"(OCTA)*, pages 1–5. IEEE.
- Fidahić, B. (2021). Case study: The International Criminal Tribunal for the former Yugoslavia's court transcripts in Bosnian/Croatian/Serbian—part 1: Needs, feasibility, and output assessment. *Genocide Studies and Prevention: An International Journal*, 15(2):37–48.
- Greene-Colozzi, E. A., Freilich, J. D., and Chermak, S. M. (2021). Developing open-source databases from online sources to study online and offline phenomena. In Anita Lavorgna et al., editors, *Researching Cybercrimes: Methodologies, Ethics, and Critical Approaches*, pages 169–190. Springer International Publishing, Cham.
- Hinton, A. L. (2012). Critical genocide studies. *Genocide Studies and Prevention: An International Journal*, 7(1):4–15.
- Hoang, L. and Schneider, J. (2018). Opportunities for computer support for systematic reviewing - a gap analysis. In Gobinda Chowdhury, et al., editors, *Transforming Digital Worlds – 13th International Conference, iConference 2018, Proceedings*, pages 367–377, Germany, March. Springer International Publishing.
- ICTY. (2016). *Infographic: ICTY Facts & Figures*. International Criminal Tribunal for the former Yugoslavia. <https://www.icty.org/node/9590>.
- Jurafsky, D. and Martin, J. (2021). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. 3rd (draft) edition. <https://web.stanford.edu/~jurafsky/slp3/>.
- Kaptein, R., Van den Broek, E. L., Koot, G., and Huis in 't Veld, M. A. (2013). Recall oriented search on the web using semantic annotations. In *Proceedings of the Sixth International Workshop on Exploiting Semantic Annotations in Information Retrieval (ESAIR '13)*, page 45–48, San Francisco, California, USA, October. Association for Computing Machinery.
- Kong, W. and Allan, J. (2016). Precision-oriented query facet extraction. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, CIKM '16*, page 1433–1442, New York, NY, USA. Association for Computing Machinery.
- Koster, C. H., Oostdijk, N. H., Verberne, S., and D'hondt, E. K. (2009). Challenges in professional search with PHASAR. In *Proceedings of the Dutch-Belgian Information Retrieval Workshop*, pages 101–102, Netherlands.
- Kruschwitz, U. and Hull, C. (2017). Searching the en-



- terprise. *Foundations and Trends in Information Retrieval*, 11(1):1–142.
- Li, Q., Peng, H., Li, J., Xia, C., Yang, R., Sun, L., Yu, P. S., and He, L. (2020). A survey on text classification: From shallow to deep learning. *arXiv preprint arXiv:2008.00364*.
- MacFarlane, A., Russell-Rose, T., and Shokrane, F. (2021). Search strategy formulation for systematic reviews: Issues, challenges and opportunities. *arXiv preprint arXiv:2112.09424*.
- Minaee, S., Kalchbrenner, N., Cambria, E., Nikzad, N., Chenaghlu, M., and Gao, J. (2021). Deep learning-based text classification: A comprehensive review. *ACM Computing Surveys*, 54(3):1–40.
- Mozafari, M., Farahbakhsh, R., and Crespi, N. (2020a). A BERT-based transfer learning approach for hate speech detection in online social media. In Hocine Cherifi, et al., editors, *Complex Networks and Their Applications VIII*, pages 928–940, Cham. Springer International Publishing.
- Mozafari, M., Farahbakhsh, R., and Crespi, N. (2020b). Hate speech detection and racial bias mitigation in social media based on BERT model. *PLOS ONE*, 15(8):e0237861.
- Müller, M., Salathé, M., and Kummervold, P. E. (2020). COVID-Twitter-BERT: A natural language processing model to analyse COVID-19 content on twitter. *arXiv preprint arXiv:2005.07503*.
- Nakov, P., Corney, D. P. A., Hasanain, M., Alam, F., Elsayed, T., Barrón-Cedeño, A., Papotti, P., Shaar, S., and Martino, G. D. S. (2021). Automated fact-checking for assisting human fact-checkers. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 4551–4558. ijcai.org.
- Noor, S. and Bashir, S. (2015). Evaluating bias in retrieval systems for recall oriented documents retrieval. *International Arab Journal of Information Technology (IAJIT)*, 12(1):53–59.
- Russell-Rose, T., Gooch, P., and Kruschwitz, U. (2021). Interactive query expansion for professional search applications. *Business Information Review*, 38(3):127–137.
- Singh, M., Jakhar, A. K., and Pandey, S. (2021). Sentiment analysis on the impact of coronavirus in social life using the BERT model. *Social Network Analysis and Mining*, 11:33.
- Sohn, H. and Lee, H. (2019). Mc-bert4hate: Hate speech detection using multi-channel BERT for different languages and translations. In *2019 International Conference on Data Mining Workshops (ICDMW)*, pages 551–559, Beijing, China, November. IEEE.
- Sun, C., Huang, L., and Qiu, X. (2019). Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentence. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 380–385, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Tong, S. and Koller, D. (2001). Support vector machine active learning with applications to text classification. *Journal of Machine Learning Research*, 2(Nov):45–66.
- Totten, S. (2017). *Advancing Genocide Studies: Personal Accounts and Insights from Scholars in the Field*. Routledge.
- Verberne, S., He, J., Wiggers, G., Russell-Rose, T., Kruschwitz, U., and de Vries, A. P. (2019). Information search in a professional context - exploring a collection of professional search tasks. *CoRR*, abs/1905.04577.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. (2019). GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR 2019*, New Orleans, Louisiana, May.
- Webersinke, N., Kraus, M., Bingler, J. A., and Leipold, M. (2021). ClimateBERT: A pretrained language model for climate-related text. *arXiv preprint arXiv:2110.12010v1*.
- Xu, H., Liu, B., Shu, L., and Yu, P. (2019). BERT post-training for review reading comprehension and aspect-based sentiment analysis. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2324–2335, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Zhang, W., Yoshida, T., and Tang, X. (2007). Text classification based on multi-word with support vector machine. In *2007 IEEE International Conference on Systems, Man and Cybernetics*, pages 3519–3524, Montreal, Canada. IEEE.
- Zhang, T., Wu, F., Katiyar, A., Weinberger, K. Q., and Artzi, Y. (2021). Revisiting few-sample BERT fine-tuning. In *International Conference on Learning Representations*, Vienna, Austria.