

Automating Violence Detection and Categorization from Ancient Texts

Alhassan Abdelhalim and Michaela Regneri

Universität Hamburg, Dept. of Computer Science, Hamburg, Germany
{alhassan.abdelhalim, michaela.regneri}@uni-hamburg.de

Abstract

Violence descriptions in literature offer valuable insights for a wide range of research in the humanities. For historians, depictions of violence are of special interest for analyzing the societal dynamics surrounding large wars and individual conflicts of influential people. Harvesting data for violence research manually is laborious and time-consuming. This study is the first one to evaluate the effectiveness of large language models (LLMs) in identifying violence in ancient texts and categorizing it across multiple dimensions. Our experiments identify LLMs as a valuable tool to scale up the accurate analysis of historical texts and show the effect of fine-tuning and data augmentation, yielding an F1-score of up to 0.93 for violence detection and 0.86 for fine-grained violence categorization.

1 Introduction

Violence has been a defining element in human history, influencing cultural values, political structures, and social norms (Frier, 1985; Raaflaub et al., 2007; Konstan, 2007). Understanding its role in shaping ancient civilizations provides valuable insights into societal evolution, power dynamics, and conflict resolution (Westbrook and Beckman, 2003; Redfield, 1994; Bizos, 2008). To analyze historical texts for information on violent events, historians have traditionally relied on manual analysis, reading, and annotating vast amounts of text. While manual annotation remains a gold standard for nuanced interpretations, time and labor required for the sheer volume of ancient texts and their linguistic complexities make this approach intractable for exhaustive collections of ancient manuscripts. The rapid growth of digital archives and historical corpora underscores the need for automated methods to assist historians in extracting information more efficiently.

Large Language Models (LLMs), such as BERT

(Devlin et al., 2019), RoBERTa (Liu et al., 2019), and GPT (Radford and Narasimhan, 2018), have successfully been applied to a wide range of classification tasks, also for scaling annotation of historical texts (Celli and Mingazov, 2024). So far, they have not been used to classify text passages denoting violent events.

Our research bridges the gap between the hermeneutical processes of historical analysis and the computational methods of natural language processing. We develop and evaluate methodologies that automate the annotation of violence in ancient texts while preserving the depth of understanding traditionally achieved through manual methods. As our gold standard, we use the manually curated ERIS database (Riess and Zerjadtke, 2015)¹, a large digital collection of violent events from ancient literature.

We first identify the violent passages contained in ERIS within their original texts using classifiers based on LLMs. Then, we further reproduce some more fine-grained annotations from ERIS, categorizing the violent passages across multiple dimensions: level of violence, contextual background, underlying motives, and long-term consequences. The results of our study show that LLMs offer a promising solution for extracting violence data. They can expedite the identification of violent events and the extraction of contextual information from ancient texts. With accurate results for a range of classification tasks around violence, LLMs can complement the expertise of historians, allowing them to focus on deeper interpretative tasks rather than the extensive and time-consuming data processing typically required.

In the following, we first give an overview of related work (Sec. 2) before we introduce our dataset and methodology (Sec. 3). We then present our results (Sec. 4) and discuss their implications (Sec. 5)

¹<https://www.ancientviolence.uni-hamburg.de>

before we conclude with a short summary and ideas for future work (Sec. 6). Code and data are provided as supplementary material ².

2 Background & Related Work

This section provides some background on violence research in history and the digital humanities (Sec. 2.1). We then introduce large language models and discuss related work concerning LLMs for classification and annotation support (Sec. 2.2).

2.1 Historical Perspectives and Data on Violence

The meaning of violence is deeply shaped by cultural context, making it a complex phenomenon to define. From a historian’s point of view, violence can be defined as *"a physical act, a process in which a human being inflicts harm on another human being via physical strength"* (Riess, 2012). Violence shaped societal values, legal systems, and social hierarchies in ancient civilizations. Interpersonal violence often reflected concepts of honor, justice, and societal expectations, as reflected in texts like *The Iliad* (Diemke et al.; Konstan, 2007). Legal codes like Hammurabi’s *Lex talionis* and Roman law institutionalized violence, balancing societal order and retributive justice (Roth, 1995; Frier, 1985).

Power dynamics frequently used violence as a tool for asserting dominance, with leaders such as Julius Caesar and Augustus consolidating power through both physical and symbolic acts of violence (Fagan, 2011; Dando-Collins, 2010). Gendered violence highlighted patriarchal structures, as myths and legal frameworks depicted male dominance and societal control (Lerner, 1986; Pomeroy, 2011). Conflict resolution in ancient texts ranged from violent duels to legal settlements and diplomatic treaties, such as the peace treaty after the Battle of Kadesh (Witham, 2020; Gagarin, 1982).

Psychological drivers of violence, such as honor, revenge, and emotional turmoil, are central to narratives like *The Iliad* and *The Oresteia*, where cycles of vengeance reflect societal norms and the transition to judicial systems (Olson, 1990; Cohen, 1986). Violence in historical accounts, such as Caesar’s assassination, also humanizes figures, exposing vulnerabilities and the socio-political landscapes of their time (Tranquillus and Graves, 1962; Allen, 2005).

²<https://osf.io/ae835/>

Analyzing violence in ancient texts enables researchers to gain insights into societal evolution (Westbrook and Beckman, 2003), comparative legal systems (Trigger, 2003; Redfield, 1994), and the foundation of modern justice (Jackson, 1968; Bizos, 2008; Eichler, 2009). Detecting violent instances in ancient texts presents unique challenges due to the implicit and symbolic nature of violence in historical narratives.

In digital humanities, the study of violence in ancient texts relies on digital resources which provide access to extensive literary and historical collections. In our work, we focus on two of these resources:

Perseus³ (Smith et al., 2000) offers Greek and Roman literature with translations, linguistic annotations, and open-access tools, enabling tasks like text reconstruction and model training (Assael et al., 2019). Despite its utility, it faces usability challenges (Lang, 2018; Preece and Zepeda, 2009).

ERIS (Riess and Zerjadtke, 2015) is a curated and expanding database of violent depictions in Greek, Roman, and some medieval texts. It includes metadata for bibliographic contexts and details of violent events. We provide a more detailed discussion of ERIS compared to Perseus because ERIS plays a central role in our study and has significant potential for future expansion. In contrast, Perseus, being a widely recognized and extensively documented resource, primarily served as a supplementary source to retrieve non-violent contexts for our dataset. ERIS is further introduced in Sec. 3.1.

2.2 Large Language Models

Large Language Models (LLMs) have revolutionized Natural Language Processing (NLP), enabling advanced text understanding and generation capabilities that were previously unattainable. Built on the architecture of Transformers (Vaswani et al., 2017), LLMs such as Generative Pre-trained Transformer (GPT) (Radford and Narasimhan, 2018), BERT (Devlin et al., 2019), and RoBERTa (Liu et al., 2019) have set new benchmarks in language modeling and processing tasks.

GPT excels in generative tasks like text completion and translation by leveraging a unidirectional architecture that predicts the next word based on prior context (Brown et al., 2020). In contrast, BERT (Bidirectional Encoder Representations from Transformers) introduced bidirectional

³<http://www.perseus.tufts.edu/>

context understanding, enabling deeper insights for tasks such as question answering and named entity recognition (Devlin et al., 2019). RoBERTa (Robustly Optimized BERT Pretraining Approach) further refined BERT’s capabilities by using larger datasets for training and optimizing various hyperparameters, which enhances performance across various benchmarks (Liu et al., 2019).

These models demonstrate the power of pretraining on vast datasets, capturing linguistic patterns and contextual nuances that generalize across diverse domains. In consequence, LLMs cemented their role as main component for scalable language processing, especially various classification tasks, such as sentiment analysis (Bang et al., 2023), text categorization (et al, 2023), and natural language inference (Honovich et al., 2022). The possibility to fine-tune such pre-trained models to small domains makes them a versatile tool also for uncommon data like ancient texts: They have already been used for scaling up annotation of historical data (Celli and Mingazov, 2024), and for hate speech detection (Mathew et al., 2021). Both tasks have goals close to our objective of extracting and categorizing violence from ancient texts. Our method is developed to scale the annotation of violent events in ancient texts, and we are also concerned with textually manifested ferocity. Our contribution extends previous approaches in that we use annotation methods for violent texts and that our data contains descriptions of violence rather than verbal assaults, as in hate speech. To the best of our knowledge, we present the first study that automatically extracts and annotates violence from historical text data.

3 Data and Experimental Setup

In this section, we explain ERIS as the basis for our experiments (Sec. 3.1), how we set up the experiments for violence detection (Sec. 3.2) and violence categorization (Sec. 3.3), and introduce the evaluation metrics used for both tasks (Sec. 3.4).

3.1 Data: The ERIS Database

ERIS (Riess and Zerjadtke, 2015) is a manually curated and continuously growing database containing depictions of violence from Greek, Roman and some medieval texts, including references to violence from Herodian, Plutarch, Tacitus, Thucydides and Xenophon. Each text passage is annotated with metadata, denoting the bibliographic contexts as well as details on the violent event. Among other

labels, it categorizes violent acts by context, motives, and social factors. It also provides metadata as timestamps and geographical coordinates, supporting advanced filtering and geospatial analysis. ERIS emphasizes sociological dimensions of violence, enabling a deeper understanding of its impacts across time and regions. Most notably, ERIS contains links to the Perseus database to match violence passages to their original texts. At the time of writing this paper, ERIS contained 3,252 entries spanning various time periods, starting from Archaic Greece in the 7th century BCE to the Salian period in the 11th century AD.

Attribute	Details
Related Conflict	Wars of Alexander The Great
Perpetrator	<i>Name:</i> Alexander III the Great <i>Age:</i> Adult <i>Activity:</i> Monarch/Ruler <i>Origin:</i> Macedonian
Victim	<i>Name:</i> Cleitus the Black <i>Age:</i> Adult <i>Direct Consequence:</i> death <i>Origin:</i> Macedonian
Third Party (Person)	<i>Name:</i> Aristophanes <i>Age:</i> Adult <i>Activity:</i> Soldier
Third Party (Group)	Friends of Alexander III <i>Origin:</i> Mixed <i>Age:</i> mixed <i>Activity:</i> commander/general
Source	Plutarch, Alexander 51.5
Year	328 B.C.
Location	Maracanda (Samarkand)
Time Period	Hellenistic Greece
Level	Interpersonal
Context	entertaining
Motivation	emotional
Weapon	Spear
Original Text	"οὕτω δὴ λαβὼν παρά τινος τῶν δορυφόρων Ἀλέξανδρος αἰχμὴν ἀπαντῶντα τὸν Κλεῖτον αὐτῷ καὶ παράγοντα τὸ πρὸ τῆς θύρας παραχύλιμα διελαύνει."
Translation	"And so, at last, Alexander seized a spear from one of his guards, met Cleitus as he was drawing aside the curtain before the door, and ran him through."
Remark	<i>perpetrator:</i> Alexander is shocked by his deed and tries to kill himself. This is mentioned in 51.6. <i>thirdperson:</i> The presence of these persons is mentioned in 51.1-4 and 51.6.

Figure 1: An entry from ERIS titled : Alexander kills Cleitus with a spear.

Figure 1 shows an example entry from ERIS. Each entry includes metadata such as title, source references, historical period, and century, as well as detailed classifications of violence level, context, motive, weapon, consequences, and method of execution. Additionally, it provides temporal and situational context, including date, season, month, and duration, along with references to the primary text sources. Some of the attributes also refer to information not contained in the text passage, here noted as *Remark*. ERIS mostly contains Greek and Roman literature, along with English translations. Our work is based on the ERIS content from biographies of Plutarch, an ancient Greek writer. We work with the English translations of the original texts.

3.2 Violence Detection

In our first experiment, we perform a binary classification task to detect instances of violence (and distinguish them from non-violent passages) in ancient texts. For classification, we compare the plain pre-trained models with fine-tuned LLMs.

As ERIS contains only violent passages, we additionally need comparable non-violent examples to train our model. To obtain those, we retrieve the context of the violent passages from ERIS by re-connecting them to their source texts. Then we train LLMs to distinguish violent from non-violent passages. As a baseline, we also use the ChatGPT-API to simulate an annotator that works with the support of ChatGPT and compare the results.

Data Pre-processing

To obtain data that we can use for training and testing, we need to amend the ERIS data with non-violent examples. Our core idea is to retrieve data from the original texts the ERIS passages were extracted from and use sentences that are not labeled in ERIS as nonviolent data. Because this requires us to have digital access to the respective original texts, we restrict this experiment to ERIS samples from Plutarch’s biographies, which have digital links to their source text in the Perseus database.

For each violent passage from ERIS, the full sections from which these excerpts were derived were retrieved. Any paragraph not explicitly marked as violent in ERIS was treated as non-violent, forming the negative examples for the dataset, resulting in a final dataset of 461 violent and 2103 non-violent texts extracted from 13 different Plutarch books. We assume that for any book that is completely

annotated for ERIS, each text part that is not contained in ERIS does only contain non-violent text. This assumption might not always hold, because annotators could have missed some passages. We discuss future assessment of this in the [Limitations](#) section.

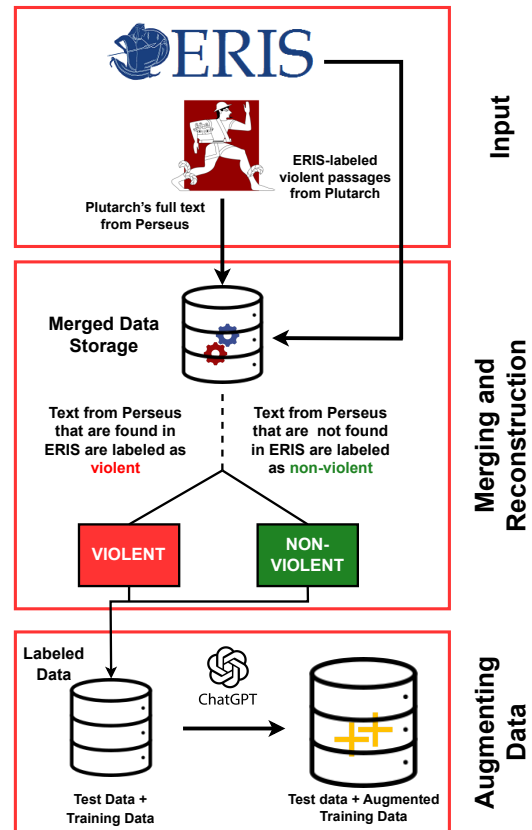


Figure 2: Data Preprocessing Pipeline for Violence Detection.

As a held-out test set for evaluation, 500 texts (371 non-violent and 129 Violent) distributed evenly across the 13 Plutarch books were selected. The remaining passages were used later on for training, fine-tuning, and data augmentation.

Because the Plutarch vitae constitute a comparably small dataset, we want to explore the impact of data augmentation with LLMs, for which we used the OpenAI API. For each entry in the training dataset, we generated three paraphrased versions and thus quadrupled the training data size. We retained 6928 non-violent texts and 1328 violent texts. Then we tested the augmented models on the 500 sentences reserved for evaluation to quantify the influence of augmentation. Figure 3 shows an exemplary sentence from ERIS and the three

paraphrases obtained using the OpenAI API.

Original Sentence	<i>As a result of this battle, the Greeks could not only harry the country of the King without fear; but had the satisfaction of seeing due punishment inflicted upon Tisaphernes, an abominable man, and most hateful to the Greek race. [Plut. Ages. 10.3]</i>
Paraphrase 1	<i>Consequently, the Greek army found themselves able to pillage the king's realm without trepidation, while also deriving satisfaction from the punishment dealt to Tisaphernes, a detestable individual who was particularly reviled by their people.</i>
Paraphrase 2	<i>This battle enabled the Greeks to wreak havoc across the king's lands unimpeded, while they also relished the sight of Tisaphernes, an odious person hated by the Greek civilization, finally facing the consequences of his actions.</i>
Paraphrase 3	<i>Following this conflict, the Hellenic forces were not only able to raid the territory of the monarch freely but also took pleasure in witnessing Tisaphernes, a loathsome figure deeply despised by the Greeks, receive just retribution.</i>

Figure 3: An example of our text augmentation. This approach effectively quadruples the training data.

Model Selection

For violence detection, we evaluated BERT-large and RoBERTa-large as LLMs for direct classification. Both models were tested first as-is and then with fine-tuning to the historical data. As an additional model that simulates annotation with the support of ChatGPT, we used the GPT-4o mini API with a specifically designed prompt that outputs the classifications. We provide the prompt in [Appendix B](#).

3.3 Categorizing Violent Events

In our second experiment, we automatically apply a more fine-grained annotation of violent texts, aiming to reproduce some ERIS annotations. In contrast to the first experiment, we use the full ERIS database as our source data. Thus, our input contains a wider variety of source texts than the violence classification (which was restricted to Plutarch biographies) and only texts that are manually labeled as violent. For this experiment, data augmentation was not suitable because we would have to augment the fine-grained annotation from ERIS as well.

With the ERIS passages, we use a multi-class classification approach across four key dimensions from the ERIS annotations:

- **Level of Violence:** Classifies instances of violence into four categories: interpersonal (conflict between individuals), intrapersonal (self-harm), intersocial (conflicts between groups, like wars), and intrasocial (conflicts within a societal group). They highlight the relational context of the events.
- **Context:** Contains 25 categories of the setting in which the violence occurred, with various political, military, and social contexts.
- **Motive:** 13 different classes for the underlying reasons for violent actions, distinguishing between tactical/strategic goals, political ambitions, adherence to authority, emotional impulses, and economic motives.
- **Long-Term Consequences:** The most fine-grained label with 38 outcomes of violent events, including social disruption, political changes, and personal impacts.

We split the dataset into 80% for training/validation and 20% for testing. Some (5) classes with very few instances do not occur in the randomly assigned test split. We fine-tuned and evaluated one BERT and one RoBERTa model per dimension.

3.4 Evaluation Metrics

For both experiments, we measure the performance of all models using the standard evaluation metrics precision, recall and F1 score.

Given that TP , FP , FN , TN are the True Positives, False Positives, False Negatives, and True Negatives respectively, key metrics are defined as follows:

$$\textbf{Precision: } P = \frac{TP}{TP + FP} \quad (1)$$

Precision measures the proportion of correct positive predictions.

$$\textbf{Recall: } R = \frac{TP}{TP + FN} \quad (2)$$

Recall measures the proportion of actual positives that are correctly identified.

$$\textbf{F1 Score: } F_1 = \frac{2 * P * R}{P + R} \quad (3)$$

The F1 score is the harmonic mean of precision and recall, which is sensitive to disparities between

them. This property ensures that the F1 score is low if either precision or recall is low, accurately reflecting the model’s overall performance.

We also provide two baselines: majority and random. A **majority baseline** represents a trivial classifier that only predicts the majority class c_{majority} . Given class probabilities $P(c_k)$ for K classes, a class c_k is predicted with: $\hat{y} = c_{\text{majority}}, \forall x \in X$.

A **random baseline** assigns labels based on class probabilities $p_i = \frac{\#C_i}{X}$. The expected probability of making a correct prediction is given by $\sum_{i=1}^N p_i^2$. This represents the probability of randomly guessing the correct label, serving as a lower-bound benchmark for classifiers.

4 Results

In this section we provide our results for violence detection (Sec. 4.1) and violence categorization (Sec. 4.2).

4.1 Violence Detection

Our results are summarized in Table 1. Overall, BERT with augmentation and fine-tuning performs best for our task. Fine-tuning enhanced the results drastically. The fine-tuned BERT and RoBERTa-large yielded an F1-score of 0.83 and 0.87, effectively capturing violent instances. Both provided competitive results.

Applying data augmentation enhanced the performance of both models. In particular, it vastly enhanced recall for all models, which is of particular interest for supporting annotators: The most common mistake when extracting violent passages manually is to miss them in the text. Having a pre-processor with high recall (maybe compromising with lower precision) can perfectly complement the precise human annotation because it is much faster to sort out falsely selected violent passages than to re-read the whole source text to retrieve missed but relevant paragraphs.

For F1, data augmentation made only a significant difference for BERT ($p < 0.05$ using McNemar’s test), but not for RoBERTa.

Our simulated zero-shot GPT annotator implemented with the general-purpose GPT-4o mini model attained an F1-score of 0.71 but struggled with non-violent instances. We attribute this to the lack of fine-tuning in ChatGPT, which is sup-

ported by both our results and many other studies that measure the importance of domain-specific fine-tuning for better classification (Rietzler et al., 2020; Rostam and Kertész, 2024; Liu et al., 2024). We also evaluated the larger GPT-4o model exploratively, which is approximately 16 times more expensive than the GPT-4o mini version. Despite the increased computational cost, GPT-4o offered only marginal improvements on our test data in the F1 score (0.5), indicating limited practical advantage for this task. We thus continued using the GPT-4o mini model.

4.2 Violence Categorization

For categorization, we used fine-tuned BERT-large and RoBERTa-large. An overview of the results is shown in Table 2. We report the averages over all instances, which amounts to weighted averages over the individual classes. A detailed breakdown by individual labels is given in Appendix D. We generally achieve promising results with an F1 score of 0.8, even for the most fine-grained category (long-term consequences with 37 classes). As for violence identification, BERT shows a slightly better performance than RoBERTa.

For identifying the *violence level*, the models performed best in classifying interpersonal and intersocial violence, achieving high precision and recall. However, intrapersonal violence posed challenges due to its lower representation and the subtle contextual understanding required.

For *context*, F1 is still comparably high given the complexity of the task with 23 classes. Looking at the details, we find that the model effectively identified broad categories like "War/Military Campaign" and "Battle" but struggled with nuanced distinctions between similar contexts, such as large-scale campaigns versus single combat.

Distinguishing *motives* works with similar accuracy. Again, the model performed well in identifying broad categories like "Tactical/Strategical" and "Political" but struggled with nuanced or less frequent categories such as "Emotional" and "Ambition". Overlaps between motives like "Political", "Following Orders", and "Tactical/Strategical" led to misclassifications.

Finding the *Long-Term Consequences* was the most challenging task with 37 different classes. The model excelled in identifying concrete categories like "Destruction/Devastation" and "Victory," which are frequently referenced in historical texts. However, categories with fewer examples,

	Model	Precision	Recall	F1-Score	Support
Violent	GPT-4o mini	0.69	0.74	0.71	129
	BERT [as-is]	0.25	0.97	0.40	129
	BERT [fine-tuned]	0.88	0.78	0.83	129
	BERT [fine-tuned and augmented]	0.87	0.99	0.93	129
	RoBERTa [as-is]	0.00	0.0	0.00	129
	RoBERTa [fine-tuned]	0.89	0.86	0.87	129
	RoBERTa [fine-tuned and augmented]	0.82	0.99	0.90	129
Non-Violent	GPT-4o mini	0.91	0.88	0.89	371
	BERT [as-is]	0.00	0.00	0.00	371
	BERT [fine-tuned]	0.93	0.96	0.94	371
	BERT [fine-tuned and augmented]	1.00	0.95	0.97	371
	RoBERTa [as-is]	0.74	1.0	0.85	371
	RoBERTa [fine-tuned]	0.95	0.96	0.96	371
	RoBERTa [fine-tuned and augmented]	1.00	0.92	0.96	371
Overall	GPT-4o mini	0.69	0.74	0.71	500
	BERT [as-is]	0.25	0.97	0.40	500
	BERT [fine-tuned]	0.88	0.78	0.83	500
	BERT [fine-tuned and augmented]	0.87	0.99	0.93	500
	RoBERTa [as-is]	0.00	0.00	0.00	500
	RoBERTa [fine-tuned]	0.89	0.86	0.87	500
	RoBERTa [fine-tuned and augmented]	0.82	0.99	0.90*	500
Baselines (overall)	Majority (all non-violent)	0.74	0.74	0.74	500
	Random	0.61	0.61	0.61	500

Table 1: Violence detection performance across different models, evaluated per class (Non-Violent and Violent). Support indicates the number of instances in each class of the test set. (*) marks an insignificant difference.

Dimension	Classes	RoBERTa			BERT			Baselines	
		Precision	Recall	F1-Score	Precision	Recall	F1-Score	Majority	Random
Level	4	0.95	0.95	0.95	0.93	0.93	0.93	0.67	0.49
Context	23	0.86	0.85	0.84	0.86	0.86	0.85	0.33	0.16
Motive	12	0.86	0.85	0.85	0.86	0.86	0.86	0.35	0.20
Consequences	37	0.82	0.80	0.80	0.82	0.81	0.81	0.36	0.16

Table 2: Overall Violence Categorization Results. Breakdowns by label are provided in [Appendix D](#)

such as "Exile" and "Coronation," proved challenging, resulting in lower precision and recall. The abstract nature of some consequences, like political changes or psychological impacts, added complexity to classification.

5 Discussion

The experiments demonstrated the potential of fine-tuned large language models (LLMs) in detecting and classifying violence in ancient texts. Our evaluation demonstrates the models' strengths in violence classification, with an F1-score of up to 0.93. In manual classification recall is often the problem due to implicit or symbolic violence, ambiguous wording, and a bias toward precision, leading to missed instances. Our finetuned and augmented models achieve a high recall, showing that LLMs can mitigate blind spots that humans miss. However, challenges like class imbalance, conceptual

overlap, and abstract categories in multi-class tasks revealed areas for improvement.

For violence categorization, our approach excelled in well-represented and concrete categories, such as "Victory" and "Destruction," but struggled with abstract or underrepresented categories like "Intrapersonal Violence" or "Exile". Conceptual overlaps, such as between "Political" and "Tactical" motives, also led to misclassifications.

From the perspective of historians, choosing between fine-tuned models, tools like ChatGPT, or manual annotation depends on specific project needs. We provide an overview over the specific features to be considered for applying fine-tuned LLMs and ChatGPT (either via user interface or via API) in [Table 3](#). Fine-tuned LLMs excel in structured, large-scale tasks where efficiency and consistency are paramount, offering rapid processing capabilities that can save months of manual labor. ChatGPT, while versatile and user-friendly,

lacks domain-specific fine-tuning, making it less reliable for specialized classifications but valuable for exploratory tasks or initial insights. Manual annotation remains irreplaceable for complex interpretative work, especially in ambiguous cases requiring deep historical expertise. A hybrid approach, where LLMs handle bulk annotation and historians validate edge cases, offers an optimal balance between efficiency and precision.

Criteria	LLM Finetuning	API
Highly specialized task	✓	✗
Requires extensive labeled data	✓	✗
Cost-effective for small tasks	✗	✓
Faster deployment	✗	✓
Full control over architecture	✓	✗
Local dependency	✓	✗
inference speed	✓	?
Suitable for dynamic scaling	✗	✓
Ongoing model maintenance	✓	✗
Scalability	✗	✓
Convenience / Usable across devices	✗	✓
Ongoing Maintenance / feedback	✗	✓
Ethical considerations	?	✓

Table 3: Pros and cons of fine-tuning LLMs vs. zero-shot approach through pre-trained OpenAI APIs

Convenience and usability are also to be considered when choosing between fine-tuning LLMs or directly using APIs. Fine-tuned models require technical expertise for setup and training but deliver streamlined workflows once operational. ChatGPT, with its accessible API and conversational interface, is more user-friendly and easy to use since it can be conveniently used in tablets or mobile phones. However, it lacks the tailored accuracy of fine-tuned models. While manual annotation is intellectually robust, it is resource-intensive and impractical for large datasets. Integrating intuitive interfaces with fine-tuned models could enhance their usability, encouraging broader adoption among non-technical users.

Inference speed varies between fine-tuned models and API-based solutions. Fine-tuned models offer lower latency but require dedicated hardware, while API-based models provide scalability but introduce network latency and rate limits. Fine-tuning is preferable for low-latency applications, while APIs offer scalability and ease of use.

Ongoing model maintenance refers to the continuous process of monitoring, updating, and retraining fine-tuned LLMs to maintain their performance and adapt to evolving data distributions or task requirements. When practitioners fine-tune their own models, they bear the responsibility for performance monitoring, infrastructure management, and regular model updates to ensure accuracy and relevance over time.

Ethical and bias considerations differ significantly between fine-tuned LLMs and API-based solutions. Pre-trained APIs are typically pre-moderated, incorporating safeguards to filter harmful or biased outputs. On the other hand, fine-tuned models require custom mitigation strategies (Jin et al., 2021; Garimella et al., 2022), which can either reduce or amplify biases, depending on dataset quality and training methods. Fine-tuning allows for domain-specific alignment but poses risks if ethical oversight is inadequate.

The implications of this research extend beyond ancient texts, offering valuable insights for analyzing contemporary violence depictions, addressing modern datasets such as media reports, social media content, or legal documents. Adapting the models to contemporary datasets would require adjustments to account for different linguistic styles, cultural contexts, and evolving definitions of violence, presenting an exciting avenue for interdisciplinary research.

A significant gap lies in automating the identification of abstract or highly contextual categories, such as psychological impacts or symbolic violence. Achieving this would require expanding datasets, understanding abstractions in LLMs (Regneri et al., 2024), incorporating knowledge bases (Wang et al., 2024), and exploring advanced techniques like retrieval augmented generation (RAG) (Chen et al., 2024). Developing dynamic models that can learn from continuous expert feedback through techniques like reinforcement learning from human feedback (RLHF) could also bridge this gap (Kaufmann et al., 2023).

6 Conclusion

In this work, we proposed a framework for automating the classification and categorization of violent ancient texts using LLMs. Our two main contributions are the development of models capable of accurately classifying violent sentences, and employing these models to automate the process of

fine-grained violence categorization. In both cases, we showed the effect of fine-tuning the models. For violence detection, we also showed that data augmentation drastically enhances recall, which is the most important measure for supporting manual annotation. Our results can enable historians to accomplish tasks that previously required months or years in minutes. To the best of our knowledge, we are also the first to utilize the OpenAI API to classify violent ancient historical texts and compare its performance against other pre-trained models. Our findings underscore the potential of LLMs to automate labor-intensive tasks and pave the way for large-scale text analysis in historical research. While fine-tuned LLMs provide structured and efficient classification, ChatGPT remains useful for exploratory tasks, and manual annotation retains its importance in complex interpretative work.

Challenges remain, particularly with underrepresented classes and computational constraints. Exploring larger models could enhance contextual understanding while maintaining runtime performance. Future work in close collaboration with historians could help resolve ambiguous cases that even human experts find difficult to classify. A hybrid approach integrating automated classification with expert validation would maximize both efficiency and accuracy. Additionally, incorporating surrounding textual context instead of analyzing passages in isolation could further enhance classification performance. Our methods also offer potential for extending the ERIS database to annotate and include texts from more recent historical periods. Adapting the models to contemporary datasets would require adjustments for linguistic style, cultural contexts, and evolving definitions of violence, presenting exciting opportunities for interdisciplinary research.

Limitations

Our study shows a promising approach to scaling up the annotation of violent events in ancient texts. While delivering accurate results in our experiments, we acknowledge several limitations rooted in the dataset, the methodology and the experimental coverage.

Dataset and annotation: ERIS is a well-curated dataset and contains the largest amount of manually annotated violent text passages from historical texts. However, this dataset also has its limits: First, for a machine learning approach the number of ex-

amples is still comparably small. Second, it only contains historical data from ancient texts as well as some medieval texts. While we assume that our approach would be applicable (possibly after more fine-tuning) to other texts, too, we cannot evaluate it with the given data. Further, ERIS does not contain information on inter-rater agreement, so we do not have a manual comparison stating how complex the task is for humans. We also do not have a detailed account on the amount of time it takes to annotate the violence passages manually. What we do know is that it strongly depends on the annotator, and that manual efforts are, overall, not easy to scale.

Methodology and Experiments: Given the limits of the database, our experiments have further limitations added. First of all, we only operate on translations rather than original texts. This might be a restriction for both text understanding and scaling the methods to texts for which no translations are available. Currently, this mirrors the manual annotation process, because annotators with knowledge of Latin or Ancient Greek are hard to find, so most of ERIS is annotated using the translations.

For violence classification, we only used the texts available in the Perseus database, because we needed to extend the ERIS data with comparable passages that do not contain violent data. Like this, the violence classification does not contain the whole ERIS database, especially not the medieval texts. While we are convinced that our results can still carry over to other epochs and text sorts, our experiments do not prove this as of yet.

Some accuracy in the fine-grained violence categorization is lost in the automated annotation, which is partly due to the ambiguity within the texts, and partly due to the challenging fine-grained taxonomy in ERIS. It is up to future work to decide whether the actual annotation guidelines and the categories need to be adapted or whether the methodology should account for this. To make this distinction, more detailed analysis and data on inter-annotator agreement would be needed (see above).

Weighted averages were chosen to reflect overall model performance effectively, particularly given the significant imbalance between class sizes. However, this method inherently favors dominant classes and can obscure weaker results in less frequent categories. A more balanced approach should be considered, potentially involving class-

based weighting or specialized metrics to ensure accurate representation across all classes.

Further, we only used four of the fine-grained ERIS categorizations for annotations. We did not do further categorization and information extraction to simulate a complete annotation of an ERIS entry. While we think that some categories are straight-forward to apply (like the identification of the weapon), others might be impossible for a model to guess, because they are not contained in the violent passages (like geographical data or sometimes the actors). In order to do this comprehensive annotation automatically, we would have to implement a different classification approach that takes the context of the violent text passages into account. We leave this experiment for future work.

Ethics Statement

We provide an experiment that helps to classify violent text passages, primarily in ancient texts. We did not use or produce any sensitive data during those experiments. We do see the potential for our method to be applied for the common good, especially when adapted to contemporary data. Like other studies on hate speech have shown, the automated detection of harmful content can support the automated analysis of the media with the aim of *protecting vulnerable groups*.

While the methodology presented in this work is primarily intended for academic and educational purposes, we recognize the potential *misuse* of AI technologies in misrepresenting historical data when applied without supervision. A misclassification of violent text or a blind reliance on the comprehensiveness of the method can lead to unwanted mistakes in the aforementioned protective purposes. Like most statements here, this applies to basically all automation methods and needs to be mediated accordingly.

Bearing in mind the general societal awareness of jobs being automatized, our work explicitly encourages the *responsible use of AI in humanities research*. Our models are designed to complement human expertise, ensuring that tedious workload is alleviated, which might be especially welcome in the case of violent texts. Like all automation approaches, this aims at scaling in terms of data set size rather than replacing analysis depth. This allows historians to focus on deeper interpretative analyses, fostering a collaborative approach between human expertise and machine learning.

Acknowledgments

The authors wish to thank Werner Rieß and Justine Diemke for providing access to the ERIS database and offering valuable feedback on results and potential research directions. Special thanks to Sri Gowry Sritharan for her assistance in extracting non-violent examples from Perseus. Additionally, the authors extend their gratitude to Sören Laue, Hanna Herasimchyk, Lennart Bengtson, and Mostafa Kotb for their technical and methodological advice and for proofreading and improving the manuscript. We also thank the three anonymous reviewers for their helpful comments. All remaining errors are, of course, our own.

References

- Brooke Allen. 2005. [Alexander the great: Or the terrible?](#) *The Hudson Review*, 58(2):220–230.
- Yannis Assael, Thea Sommerschild, and Jonathan Prag. 2019. [Restoring ancient text using deep learning: a case study on Greek epigraphy.](#) 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 6368–6375, Hong Kong, China. Association for Computational Linguistics.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. [A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity.](#) In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718, Nusa Dua, Bali. Association for Computational Linguistics.
- George Bizos. 2008. [Ethics, politics and law in ancient greece and contemporary south africa.](#) *Phronimon*, 9(2):5–15.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners.](#) In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

- Fabio Celli and Dmitry Mingazov. 2024. [Knowledge extraction from llms for scalable historical data annotation](#). *Electronics*, 13(24).
- Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024. [Benchmarking large language models in retrieval-augmented generation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17754–17762.
- David Cohen. 1986. [The theodicy of aeschylus: Justice and tyranny in the oresteia](#). *Greece and Rome*, 33(2):129–141.
- Stephen Dando-Collins. 2010. *The Ides: Caesar’s Murder and the War for Rome*. Turner Publishing Company.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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. Association for Computational Linguistics.
- Justine Diemke, Felix K Maier, Superhero Comics, Eleonora Sereni, Ulrich Bröckling, Barbara Korte, and Ulrike Zimmermann. [helden. heroes. héros](#).
- Barry L. Eichler. 2009. [Law and morality in ancient near eastern thought](#). In *Ethics, Politics, and Democracy: From Primordial Principles to Prospective Practices*. The MIT Press.
- Liang et al. 2023. [Holistic evaluation of language models](#). *Trans. Mach. Learn. Res.*, 2023.
- Garrett G. Fagan. 2011. *The lure of the arena : social psychology and the crowd at the Roman games*. Cambridge University Press.
- Bruce W. Frier. 1985. *The rise of the Roman jurists : studies in Cicero’s Pro Caecina*. UT Back-in-Print Service.
- Michael Gagarin. 1982. [The organization of the gortyn law code](#). *Greek, Roman, and Byzantine Studies*, 23(2):129–146.
- Aparna Garimella, Rada Mihalcea, and Akhshay Amarnath. 2022. [Demographic-aware language model fine-tuning as a bias mitigation technique](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 311–319, Online only. Association for Computational Linguistics.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansky, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. [TRUE: Re-evaluating factual consistency evaluation](#).
- Bernard S Jackson. 1968. Evolution and foreign influence in ancient law. *The American Journal of Comparative Law*, 16(3):372–390.
- Xisen Jin, Francesco Barbieri, Brendan Kennedy, Aida Mostafazadeh Davani, Leonardo Neves, and Xiang Ren. 2021. [On transferability of bias mitigation effects in language model fine-tuning](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3770–3783, Online. Association for Computational Linguistics.
- Timo Kaufmann, Sarah Ball, Jacob Beck, Eyke Hüllermeier, and Frauke Kreuter. 2023. [On the challenges and practices of reinforcement learning from real human feedback](#). In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 276–294. Springer.
- David Konstan. 2007. *The emotions of the ancient Greeks: Studies in Aristotle and classical literature*. University of Toronto Press.
- Sarah Lang. 2018. [Perseus digital library](#). *RIDE: A Review Journal for Digital Editions and Resources*, 8.
- Gerda Lerner. 1986. *The creation of patriarchy*. Oxford University Press.
- An Liu, Zonghan Yang, Zhenhe Zhang, Qingyuan Hu, Peng Li, Ming Yan, Ji Zhang, Fei Huang, and Yang Liu. 2024. [PANDA: Preference adaptation for enhancing domain-specific abilities of LLMs](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 10960–10977, Bangkok, Thailand. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv*, abs/1907.11692.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2021. [Hatexplain: A benchmark dataset for explainable hate speech detection](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17):14867–14875.
- S. Douglas Olson. 1990. [The stories of agamemnon in homer’s odyssey](#). *Transactions of the American Philological Association (1974-)*, 120:57–71.
- Sarah Pomeroy. 2011. *Goddesses, whores, wives, and slaves: Women in classical antiquity*. Schocken.
- Emily Preece and Christine Zepeda. 2009. [The perseus digital library: A case study](#). *Texas ScholarWorks*.
- Kurt A Raaflaub et al. 2007. *War and peace in the ancient world*. Wiley Online Library.

- Alec Radford and Karthik Narasimhan. 2018. [Improving language understanding by generative pre-training](#).
- James M Redfield. 1994. *Nature and Culture in the Iliad: the Tragedy of Hector*. Duke University Press.
- Michaela Regneri, Alhassan Abdelhalim, and Soeren Laue. 2024. [Detecting conceptual abstraction in LLMs](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4697–4704, Torino, Italia. ELRA and ICCL.
- Werner Riess. 2012. *Performing Interpersonal Violence Court, Curse, and Comedy in Fourth-Century BCE Athens*. de Gruyter.
- Werner Riess and Michael Zerjadtke. 2015. [Eris: Hamburg information system on greek and roman violence](#). *Digital Classics Online*, pages 70–75.
- Alexander Rietzler, Sebastian Stabinger, Paul Opitz, and Stefan Engl. 2020. [Adapt or get left behind: Domain adaptation through BERT language model fine-tuning for aspect-target sentiment classification](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4933–4941, Marseille, France. European Language Resources Association.
- Zhyar Rzgar K Rostam and Gábor Kertész. 2024. [Fine-tuning large language models for scientific text classification: A comparative study](#). In *2024 IEEE 6th International Symposium on Logistics and Industrial Informatics (LINDI)*, pages 000233–000238. IEEE.
- Martha T Roth. 1995. *Law collections from Mesopotamia and Asia minor*. Scholars Press.
- David A Smith, Jeffrey A Rydberg-Cox, and Gregory R Crane. 2000. [The perseus project: A digital library for the humanities](#). *Literary and Linguistic Computing*, 15(1):15–25.
- Caius Suetonius Tranquillus and Robert Graves. 1962. *The Twelve Caesars... Translated by Robert Graves*. Cassell; printed in Czechoslovakia.
- Bruce G Trigger. 2003. *Understanding early civilizations: a comparative study*. Cambridge University Press.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Xi Wang, Liana Mikaelyan, Taketomo Isazawa, and James Hensman. 2024. [Kblam: Knowledge base augmented language model](#). *arXiv preprint arXiv:2410.10450*.
- R. Westbrook and G.M. Beckman. 2003. *A History of Ancient Near Eastern Law*. Number Bd. 2 in A History of Ancient Near Eastern Law. Brill.
- Dorothy Natalie Witham. 2020. *The battle of kadesh: Its causes and consequences*. *Master Of Arts, University Of South Africa*.

A Code and Data

We provide the training data for both tasks as well as the code, downloadable under

<https://osf.io/ae835/>

Violence detection

The folder `1_violence_detection` contains the training and test data for the violence identification task. The sentences are a subset of ERIS extended with their original contexts extracted from the Perseus database. We provide both the original dataset and the augmented dataset used for training. The annotation contains the source as noted in Perseus (book, chapter, and section), the passage text, and the violence annotation (1 for violent, 0 for non-violent).

Violence categorization

The folder `2_violence_categorization` contains a condensed version of the ERIS database, including the text passage with the four annotation dimensions we used for classification. To reproduce our training and test data, please use the code we provide.

Code

We provide two Jupyter notebooks (`violence_detection.ipynb` and `violence_categorization.ipynb`) to reproduce our data preprocessing, model training, and evaluation for both tasks.

B GPT-4o mini Testing Prompt

You are a historian that classifies historical texts into violent or non-violent based on the provided examples. The following principles apply to the classification of violent acts:

- *Arrests of people and banishments are initially recorded as acts of violence and discussed with the team before being activated.*
- *Fictional narratives, such as the conquest of Troy, are included.*
- *Establishment of colonies, verbal violence (insults), and damage to property (including fires in buildings, etc.) are excluded.*

Your task is to classify each passage based on the criteria above. Respond with only [VIOLENT] or [NON-VIOLENT] for each classification.

C GPT-4o mini Augmentation Prompt

You are a historian that wants to paraphrase sentences to create new ones for enhancing your dataset. Generate three different ways to rewrite the following sentence while keeping the same meaning. Important to note that you are not allowed to change context, motive or consequences.

D Detailed Breakdown for the Violence Categorization results

These are the extended results for [Table 2](#).

	RoBERTa Results				BERT Results			
	Precision	Recall	F1-Score	Support	Precision	Recall	F1-Score	Support
Interpersonal	0.92	0.91	0.91	96	0.93	0.88	0.90	96
Intrasocial	0.95	0.83	0.89	72	0.95	0.78	0.85	72
Intersocial	0.96	0.98	0.97	371	0.94	0.99	0.96	371
Intrapersonal	0.84	0.94	0.89	17	0.76	0.76	0.76	17
Overall	0.95	0.95	0.95	556	0.93	0.93	0.93	556

Baselines: Majority = 0.67, Random = 0.49

Table 4: Comparison of Level Results for RoBERTa, BERT, and Baselines

	RoBERTa Results				BERT Results			
	Precision	Recall	F1-Score	Support	Precision	Recall	F1-Score	Support
Civilian	1.00	0.69	0.82	29	0.96	0.79	0.87	29
Jurisdictional	0.86	0.80	0.83	30	1.00	0.77	0.87	30
War/Military Campaign	0.80	0.94	0.87	181	0.83	0.97	0.89	181
Battle	0.93	0.81	0.87	69	0.92	0.88	0.90	69
Plunder	0.69	0.53	0.60	17	0.75	0.53	0.62	17
Ambush	0.85	0.73	0.79	15	1.00	0.67	0.80	15
Conspiracy	0.82	0.82	0.82	11	0.53	0.82	0.64	11
Revolt	1.00	1.00	1.00	21	1.00	1.00	1.00	21
Conquest	0.50	0.57	0.53	7	0.57	0.57	0.57	7
Naval Battle	1.00	1.00	1.00	2	0.29	1.00	0.44	2
Religious	1.00	1.00	1.00	6	0.67	0.33	0.44	6
Institutional	0.60	0.75	0.67	4	1.00	0.75	0.86	4
Sack	0.00	0.00	0.00	1	0.00	0.00	0.00	1
Single Combat	1.00	0.50	0.67	4	1.00	0.50	0.67	4
Siege	0.83	0.81	0.82	31	0.89	0.81	0.85	31
Unknown	1.00	1.00	1.00	5	1.00	0.80	0.89	5
Regicide	0.69	1.00	0.81	11	0.79	1.00	0.88	11
Military	0.90	0.87	0.89	93	0.91	0.90	0.91	93
Entertaining	0.60	0.43	0.50	7	0.60	0.43	0.50	7
Mutiny	1.00	0.75	0.86	8	1.00	0.75	0.86	8
Familicide	1.00	1.00	1.00	2	0.00	0.00	0.00	2
Fratricide	0.00	0.00	0.00	1	0.00	0.00	0.00	1
Paramilitary	1.00	1.00	1.00	1	0.00	0.00	0.00	1
Overall	0.86	0.85	0.84	556	0.86	0.86	0.85	556
Baselines: Majority = 0.33, Random = 0.16								

Table 5: Comparison of Context Results for RoBERTa, BERT, and Baselines

	RoBERTa Results				BERT Results			
	Precision	Recall	F1-Score	Support	Precision	Recall	F1-Score	Support
Unknown	1.00	0.80	0.89	20	0.81	0.65	0.72	20
Political	0.84	0.86	0.85	122	0.91	0.86	0.89	122
Tactical/Strategical	0.87	0.88	0.87	197	0.92	0.88	0.90	197
Economical	0.74	0.82	0.78	28	0.69	0.86	0.76	28
Following Orders	0.90	0.86	0.88	77	0.81	0.90	0.85	77
Self-Defence	0.75	0.69	0.72	13	0.73	0.62	0.67	13
Emotional	0.97	0.77	0.86	43	0.92	0.84	0.88	43
Ambition	0.71	0.83	0.76	35	0.64	0.83	0.72	35
Social	0.71	1.00	0.83	5	1.00	1.00	1.00	5
Religious	0.83	0.83	0.83	6	0.83	0.83	0.83	6
Other	1.00	1.00	1.00	6	1.00	0.83	0.91	6
None/Accident	0.75	0.75	0.75	4	0.75	0.75	0.75	4
Overall	0.86	0.85	0.85	556	0.86	0.86	0.86	556
Baselines: Majority = 0.35, Random = 0.20								

Table 6: Comparison of Motive Results for RoBERTa, BERT, and Baselines

	RoBERTa Results				BERT Results			
	Precision	Recall	F1-Score	Support	Precision	Recall	F1-Score	Support
Unknown	0.78	0.89	0.83	199	0.83	0.91	0.87	199
Campaign	0.81	0.87	0.85	28	0.82	0.82	0.82	28
Conquest	0.83	0.83	0.83	24	0.58	0.92	0.71	24
Coronation/Inauguration	1.00	0.67	0.80	12	0.90	0.75	0.82	12
Exile	1.00	0.67	0.80	6	0.86	1.00	0.92	6
Death	0.81	0.72	0.72	32	0.77	0.69	0.73	54
Other	0.72	0.72	0.72	32	0.86	0.78	0.82	32
Victory	1.00	1.00	1.00	16	0.88	0.94	0.91	16
Bestowing of Honors	0.67	0.33	0.44	6	1.00	0.17	0.29	6
Issuing of Law/Decrees	1.00	0.33	0.50	3	0.50	0.33	0.40	3
Injury	0.71	1.00	0.83	5	1.00	1.00	1.00	5
Battle	0.80	0.53	0.64	15	0.67	0.67	0.67	15
Declaration of War	1.00	1.00	1.00	2	1.00	1.00	1.00	2
Retreat	0.67	0.80	0.73	10	0.67	0.80	0.73	10
Mutiny	1.00	0.50	0.67	2	1.00	0.50	0.67	2
Sending of Envoys	0.93	1.00	0.96	13	0.92	0.92	0.92	13
Civil Conflict/Civil War	0.00	0.00	0.00	1	0.00	0.00	0.00	1
Tyranny	0.50	1.00	0.67	2	1.00	1.00	1.00	2
Capture	0.71	0.71	0.71	14	0.77	0.71	0.74	14
Destruction/Devastation	0.84	0.81	0.82	26	0.84	0.81	0.82	26
Repopulation	1.00	1.00	1.00	2	1.00	1.00	1.00	2
Declaration of Peace/Truce	1.00	0.44	0.62	9	1.00	0.44	0.62	9
Release of Prisoners	1.00	1.00	1.00	2	0.67	1.00	0.80	2
Garrisoning of Troops	1.00	0.67	0.80	6	1.00	0.67	0.80	6
Famine	1.00	1.00	1.00	1	1.00	1.00	1.00	1
Siege	0.95	0.70	0.81	30	0.95	0.70	0.81	30
Deportation	1.00	0.25	0.40	4	1.00	0.50	0.67	4
Treaty/Agreement/Pact	1.00	0.33	0.50	3	0.00	0.00	0.00	3
Surrender	0.67	1.00	0.80	2	0.67	1.00	0.80	2
Financial Reward	0.75	1.00	0.86	3	0.75	1.00	0.86	3
Seclusion	0.33	1.00	0.50	2	1.00	1.00	1.00	2
Plunder	0.86	1.00	0.92	6	1.00	1.00	1.00	6
Mutilation	1.00	1.00	1.00	1	1.00	1.00	1.00	1
Revenge	1.00	1.00	1.00	6	1.00	1.00	1.00	6
Execution	0.40	0.50	0.44	4	0.33	0.25	0.29	4
Torture	0.75	1.00	0.86	3	0.75	1.00	0.86	3
Applause	1.00	0.50	0.67	2	1.00	0.50	0.67	2
Overall	0.82	0.80	0.80	556	0.82	0.81	0.81	556
Baselines: Majority = 0.36, Random = 0.16								

Table 7: Comparison of Long-Term Consequences Results for RoBERTa, BERT, and Baselines