

# A multi-level multi-label text classification dataset of 19<sup>th</sup> century Ottoman and Russian literary and critical texts

Gokcen Gokceoglu<sup>1</sup>, Devrim Cavusoglu<sup>1</sup>, Emre Akbas<sup>1</sup>, Ozen Nergis Dolcerocca<sup>2</sup>

<sup>1</sup>Department of Computer Engineering, Middle East Technical University

<sup>2</sup>Department of Modern Languages, Literatures, and Cultures, University of Bologna

## Abstract

This paper introduces a multi-level, multi-label text classification dataset comprising over 3000 documents. The dataset features literary and critical texts from 19th-century Ottoman Turkish and Russian. It is the first study to apply large language models (LLMs) to this dataset, sourced from prominent literary periodicals of the era. The texts have been meticulously organized and labeled. This was done according to a taxonomic framework that takes into account both their structural and semantic attributes. Articles are categorized and tagged with bibliometric metadata by human experts. We present baseline classification results using a classical bag-of-words (BoW) naive Bayes model and three modern LLMs: multilingual BERT, Falcon, and Llama-v2. We found that in certain cases, Bag of Words (BoW) outperforms Large Language Models (LLMs), emphasizing the need for additional research, especially in low-resource language settings. This dataset is expected to be a valuable resource for researchers in natural language processing and machine learning, especially for historical and low-resource languages. The dataset is publicly available<sup>1</sup>.

## 1 Introduction

In recent years, there has been remarkable progress in natural language processing with the introduction of large language models (Vaswani, 2017; Radford, 2018; Devlin, 2018), and well-curated, massive pre-training datasets. However, this progress has disproportionately benefited high-resource languages. For low-resource languages, the effectiveness of large language models is compromised by a number of issues including data scarcity that hinders generalization performance; tokenization processes that result

in an inadequate representation of word meaning, and the bias in digital texts toward particular topics. To address these issues, there has been an increased focus on developing more inclusive and comprehensive models, introducing diverse datasets, and establishing benchmarks.

Low-resource languages, often comprising ancient or historical linguistic forms, are a key area of inquiry in computational linguistics characterized by linguistic diversity and evolutionary trajectories over time. However, the archival preservation of written documents poses a notable challenge to computational efforts, particularly in terms of resource availability. This challenge is even present in the case of relatively recent historical languages from the 19<sup>th</sup> century, where not only limited textual resources but also non-standardized writing systems hinder computational analysis and linguistic research efforts.

With the recent advancements in deep learning research, the effort to digitize and apply machine comprehension on these type of archival resources has become a growing field of inquiry, fostering inter-disciplinary collaborations between humanities, social sciences, and computer science. One example of such endeavor that aims to overcome the limitation of lack of resources, is a large-scale project focusing on aggregation, refinement, and digitizing historical newspapers for The European Library<sup>2</sup> and Europeana<sup>3</sup> funded by the European Commission (Pekárek and Willems, 2012). In the project, several refinement methods were used on the raw content including, but not limited to Optical Character Recognition (OCR), Optical Layout Recognition, and Named Entity Recognition. The project included many languages and contributions from many national libraries

<sup>1</sup><https://huggingface.co/nonwestlit>. The dataset, code, and trained models are released under Apache 2.0 License.

<sup>2</sup><https://www.theeuropeanlibrary.org/>

<sup>3</sup><http://www.europeana.eu/>

Title	Mukaddime
Article	<p>سرمايه سعادت دارين او لان علوم وفنون نافعه تكم ممالك محروسه شاهانه ده نشر و تعميمه سعي ايلك نيت خير به سيله مقصده بعض ده و ارباب معلومانن مركب اوليق اوزره در سعادتده (جمعيت علميه عثمانيه) عو ايله بر جمعيت تشكك ايتش و بودخي سلطنت سنيه تكم تربيه عامه حلقه سزاو اربوريلان مساعي جيله سنيه موافق بولنديجي جهته نژد عاليده رهن تحصيل اوله رقي اولبايد. اراده مكار حضرت پادشاهي شرفلق بيورلمشيدى. اشو جمعيت بعض محال موقتده در فاج ماه عقد مجلس ايد رك نظمات لازمه سني وضع و تاسيس ايله بعض اثار نافعه تهينه سنيه چاشقده بوللمش ايسه ده اكثريا بو مقوله تشبثات خيرييه داييت حاده دوجار انواع مشكلات اولديغندن دائمي صورنده مركز جمعيت اتخاذ اولنلق اوزره هنوز بر محل مناسب تداركده مسترس اولنه مديجي جهته شمدي به قدر بروج مقرر اجراي مقاصدينه شروع ايلك ميسر اوله ماشيدى * مجد دينيان سلطنت و ناشر اواز بر و عاقلقت اولان ذات فضائلست حضرت پادشاهي اجداد كرامتي تخت مسعاليخته شوكتبخش جلوس اولنقاري روز روز فيروزدينرو صنوف تبعه وز بردستانلريك هر وجهه انتظار و سعادت حالري وسائلك استحصالي اخس امال جناب شهر ياري اوله رقي بجمده تعالى مغبوط اصعار سلف اولان ايام سعادت اتمام سلطنتاري حقاووز بازار ترقى و عرفان ايكن جمعيت مذكور هنگ شو حاده فائسي كرك اعضاسلك وكرك حصول امانتي ارزو اين ساتر ذوات حميت سماتك زياده موجب تاسفي اوله رقي قدر تاري مساعد اولديني مرتبه فداكارلق اختياريله عهد قريده لازم كالن محل مخصوص تداركي اسبابك استحصاله غيرت و حسيما بسا عد التقدير خدمات منصوره سنگ ايفاسنه بخلوص ابال مياشتريت ايسكه بوكره عزيمت ايتمش اولرليه جمله وظائفندن بولنان اشو ) مجموعه فنون ( عوايله معنون ائرينك دخي سنيه جديده ابتدا سندن بدا ايله ثورينه مياشت ايلماردر</p>
Categories	Cultural practices, Education, Institutions

(a) An Ottoman example article.

Title	Puteshestviia i otkrytiia doktora Barta v Afrike. Reisen und Entdeckungen in Nord- und Centralafrika in den Jahren 1849—55. Von Dr. Heinrich Barth. Tagebuch seiner im Auftrage der britischen Regierung unternommenen Reise. Bd. 1 und 2. Gotha, 1857. (Prodolzhenie.)
Article	<p>Не смотря на разнообразие требовавший и стремлений нашей современной критики изящного, можно, кажется, усмотреть два основные начала в ее оценке текущих произведений словесности. Начала эти и прежде составляли предмета, деятельной полемики между литераторами, а в последние пять лет, они обратились почти в единственный серьезный вопрос, возникавший от времени до времени на шумном поле так называемых обозреши, заметок, журналистики. Начала, о которых говорим, по существу своему еще так важны в отношении к отечественной литературе, еще так исполнены жизни и значения для нее, что там только и являлось дельное слово, где они были затронуты, там только и пропадали личные страсти и легкая работа присяжного браковщика литературы, где они706 русскш вьстннкъ. выступали на первый планъ. Спешимъ сказать, что начала эти были: идея о художественности произведен* и идея о народности, Посильное определеше ихъ съ одной стороны, а съ другой стороны сомньше въ возможности и польза; опредлешя ихъ составляли и составляютъ доселе спорный пунктъ</p>
Categories	Cultural practices, Institutions

(b) A Russian example article.

Figure 1: A training instance from Ottoman (a) and Russian (b) collection samples from third level (i.e. Cultural Discourse → Modernization Subject) . The articles are truncated for better visual appearance.

having over 16 million items. It largely covered main Western European languages. There were, therefore, several limitations for those languages that were out of the scope of the project, which resulted in limitations for refinement techniques for those particular non-European languages (e.g. Ottoman) (Neudecker and Wilms, 2013).

In order to address the limitations of digitization and refinement of non-European languages, this paper focuses on historical non-Western languages, primarily Ottoman and Russian literature texts from the 1830s to the 1910s, the most prolific years for periodical publications in the long 19<sup>th</sup> century. The rise of mass readership in periodical press roughly corresponds to these dates across all cultures with mass printing presses and (proto) print capitalism (Vincent, 2000). While Russian is not considered a low-resource language, computational efforts, particularly for nineteenth-century texts, remain limited (Martinsen, 1997). Ottoman Turkish, on the other hand, presents a unique case. It was a written language distinct from the vernacular, featuring vocabulary and syntactic elements borrowed from Persian and Arabic (Kabacalı, 2000). The orthography was also not standardized, being written in the Arabo-Persian alphabet with variations in spelling and punctuation. This is in a stark contrast to modern Turkish which has undergone systematic phonocentric standardization toward the vernacular and is written in the Latin alphabet. Therefore, any computational study conducted in modern Turkish would have significant limitations in Ottoman Turkish, as in any non-standardized language

(Ryskina, 2022).

In this paper, we provide an open-access, multi-label text classification dataset consisting of Ottoman and Russian literary texts from the 19<sup>th</sup> century. The dataset collection was prepared with due diligence from data collection to labeling. We used an OCR pipeline and developed a web-based platform for a team of experts to label the digitized content. Our dataset collection has a specific hierarchical structure (see Appendix A) and is annotated according to pre-determined categories as either single-label or multi-label. Example instances from both the Ottoman and Russian datasets are given in Figure 1.

We expect that these text classification datasets will prove instrumental in advancing future research endeavors. The availability of these datasets will enable researchers to train a wide array of models tailored to specific applications. This includes the automated categorization of digitized resources, offering enhanced efficiency and accuracy in managing digital information. We also note that the Russian and Ottoman datasets consist mainly of literary and critical texts, which offer a potential for unsupervised learning.

## 2 Related Work

A key milestone behind the recent breakthroughs in NLP is the extensive pre-training datasets. These datasets are collected from digital sources, where Western languages, predominantly English, are the primary languages represented. Specifically, by January 2024, W3Techs estimates that Western languages make up 83.1% and English 51.7%, of

the websites (W3Techs, 2024).

In recent years, efforts have been made to bridge this gap by introducing of multilingual and low-resource datasets and developing methods to improve model performance. These studies have targeted different NLP tasks, including machine translation (Guzmán et al., 2019; Mueller et al., 2020), text classification (Cruz and Cheng, 2020; Zaikis and Vlahavas, 2023; Fesseha et al., 2021; Liebeskind and Liebeskind, 2020), and part of speech tagging (Şahin and Steedman, 2019). Efforts have been made to develop new augmentation techniques (Şahin and Steedman, 2019) and training strategies (Tang et al., 2018).

Gopidi and Alam (2019) proposed an English dataset spanning the late 19<sup>th</sup> to late 20<sup>th</sup> centuries to examine the shift between poetry and prose resources. Cruz and Cheng (2020) introduced two datasets containing over 14,000 samples in Filipino for binary and multi-label classification tasks. Guzmán et al. (2019) proposed datasets for Nepali-English and Sinhala-English translations, which provide benchmarks for evaluating methods trained on low-resource language pairs. Mueller et al. (2020) used the Bible for machine translation tasks in 1107 languages, creating a multilingual corpus by varying the number and relatedness of source languages. Fesseha et al. (2021) provided a dataset curated for multi-label classification and a pre-training dataset in Tigrinya, a Semitic language spoken primarily in Eritrea and northern Ethiopia. Regatte et al. (2020) introduced a dataset for sentiment analysis in Telugu, which is one of the most spoken languages in India, but lacks adequate digital resources. Bansal et al. (2021) presented a dataset in Sumerian, one of the earliest known written languages, and proposed a cross-lingual information extraction pipeline. As a low-resource language that is no longer in use, Ottoman Turkish has been minimally used in NLP research. One such study used traditional machine learning methods to classify poems from the 15<sup>th</sup> to 19<sup>th</sup> centuries by period and genre (Can et al., 2012, 2013). Besides these low-resource dataset proposals and models trained on top of them, there are studies on improving training/fine-tuning approaches for low-resource languages in a multilingual setting (Lankford et al., 2023; Ogueji et al., 2021).

In this paper, we present the first expert-curated and labeled dataset in Ottoman Turkish and Russian and propose the first study employing

LLMs for multi-level classification tasks in both languages.

### 3 Data Collection

Our dataset collection procedure had two steps: article curation and labeling by human experts.

#### 3.1 Article Curation

Our main resources are hard copies and soft copies of texts from literary periodicals from the 1830s to 1910s, gathered from main archival resources including national and imperial libraries and digitized corpus. Experts determined a range of significant literary journals from the period, paying special attention to the wide representation of the ideological and cultural spectrum of the period. The hard or digitized copies of the selected issues of these journals were then processed and segmented into sections for each entry in a particular issue.

Due to the predominant presence of hard copy resources, particularly in Ottoman Turkish, it was necessary to digitize some of these documents to make them processable by computers. To this end, we have developed a rigorous digitization protocol aimed at transforming physical articles into digital content. This process is structured into two sequential stages: OCR followed by correction and refinement of the digitized text to ensure fidelity to the original content. OCR process for documents in Ottoman Turkish is conducted through Google Lens<sup>4</sup>, as the other OCR programs we experimented with, Abbyy (abb) and Tesseract (Arabic / Persian) (Kay, 2007), yielded low accuracy. Each PDF document page is converted to images, uploaded to Google Photos automatically via Google Cloud, and digitized by archive workers. The polishing and corrections on the digitized documents were conducted and completed by the team of experts also designated for the labeling task. We chose not to perform a formal evaluation due to time constraints and resource limitations, but the measures implemented have provided us with a high degree of confidence in the transcription accuracy.

We built a system with a web-based user interface where users can interact and label the digitized documents accordingly. The designated team of experts was signed up to the system, where each user can interact with a document reader and

<sup>4</sup><https://lens.google/>

is able to view the original images of the same document that are fed into the OCR pipeline. This was done to reduce errors in the correction phase of the digitized text.

### 3.2 Labeling

As reading, understanding and actively interpreting these texts are required for their annotation, we built a team of experts for each area, based on their academic background, linguistic skills, and scientific output regarding the nineteenth-century literary culture. The team of experts is assigned tasks to segment and label these digitized articles to form a text classification dataset. The annotations underwent a randomized cross-check process by the designated expert-leader of each language. The experts also tagged documents with bibliometric metadata in addition to labeling the determined category of that article type. This is crucial for the digitization process such that, the dataset and the database we built in parallel could be searchable in a more comprehensive way with bibliometric information which we believe will ease the search process for researchers.

The taxonomy represents the formal and thematic elements in the cultural and aesthetic world of the period. Texts are first categorized based on their form, i.e. poetry, critical article, or short story. Critical articles, labeled “cultural discourse,” are then categorized according to the subject of their content.

## 4 Dataset Analysis

Our dataset comprises articles from Russian and Ottoman sources spanning the late 19<sup>th</sup> and early 20<sup>th</sup> centuries. The structures of datasets have a hierarchical scheme, where the datasets follow 4 levels of categorization. The articles are primarily grouped into 3 categories based on their relationship to literary and cultural phenomena, namely “Literary Text,” “Cultural Discourse” and “Other,” which marks the first level of categorization (L1). “Other” refers to everything that remains outside of literary or cultural texts, such as news articles, advertisements, and obituaries. Literary Text and Cultural Discourse categories also have lower-level categories. There is no lower-level category for the articles in the class “Other.” The article counts and the names of the subcategories are shown in Table 1.

The categorization goes up to the fourth level of

Category	Ottoman	Russian
<b>Literary Text</b>	481	357
└ Short Story	99	44
└ Poetry	278	244
└ Translated Text	4	3
└ Novel	57	33
└ Play	8	15
<b>Cultural Discourse</b>	929	475
<i>Types</i>		
└ Article	796	272
└ Review	91	142
└ Biography	74	37
└ Letter	77	65
└ Manifesto	30	7
└ Travel writing	11	14
<i>Subjects</i>		
└ <b>Literature</b>	489	173
└ <b>Philosophy</b>	574	236
└ Politics	105	7
└ Translation	53	16
└ <b>Modernization</b>	152	132
└ <b>Identity</b>	136	32
└ <b>Language</b>	169	28
<b>Other</b>	409	226
<b>Total</b>	1819	1058

Table 1: Taxonomy and tree structure for the Ottoman and Russian dataset for first and second level categories. Subcategories having lower-level categories are marked in bold.

categories, and the tree structure displaying these lower levels is given in Table 2.

Within the Ottoman dataset, we have a collection of 1,819 articles from 685 journals, with an average of 1,005.29 words, and 43.4 sentences per article. The Russian dataset contains 1,058 articles from 198 journals, with an average of 4,630.69 words and 212.26 sentences per article.

## 5 Modeling experiments

We conducted multiple experiments using Ottoman and Russian datasets to establish baseline results for classification tasks.

### 5.1 Experimental Setup

**Modeling Setup.** The selection of appropriate language models is crucial for creating an effective experimental setup. Factors such as multilingual support, extensive pre-training corpora, and models including billion-parameter LLMs (Ouyang et al., 2022; Almazrouei et al., 2023; Touvron et al., 2023a) should be considered. Given these considerations, we chose two recent open-source LLMs with large and preferably multilingual pretraining data, namely Llama-2 (Touvron et al.,



Category	Ottoman	Russian
<b>Literature Genre</b>	490	173
└ Poetry	152	55
└ Theater	37	48
└ Prose Fiction	86	36
<b>Literary Movement</b>	490	173
└ Decadence	20	1
└ Symbolism	28	-
└ Romanticism	70	21
└ Traditionalism	33	-
└ Realism	49	20
└ Classicism	36	16
└ Naturalism	23	5
└ Modernism	52	-
└ Sentimentalism	2	8
<b>Philosophy</b>	576	236
└ Political Philosophy	202	173
└ State	57	70
└ Race	25	4
└ Empire/Colony	18	16
└ Nation/Society	133	132
└ Economy	46	63
└ Law	32	52
└ Class/Capital	12	38
└ Epistemology/Ontology	113	-
└ Ethics	174	35
└ Religion/Secularism	75	20
└ Morality	112	23
└ Aesthetics	105	32
└ Didacticism	56	21
└ Aestheticism	49	29
└ Philosophy of History	152	39
└ Movement	229	25
└ Enlightenment	49	4
└ Nationalism	120	7
└ Materialism	31	6
└ Woman's Question	25	6
└ Idealism	26	8
└ Orientalism	15	-
└ Marxism	4	6
<b>Modernization</b>	152	132
└ Cultural practices	42	36
└ Education	70	56
└ Print culture	35	34
└ Institutions	39	19
└ Urbanization	22	7
<b>Identity</b>	136	32
└ Localism	102	28
└ Westernizer	65	22
<b>Language</b>	169	28
└ Language reforms	62	1
└ Stylistics	55	12
└ Linguistics	77	19

Table 2: Taxonomy for the Ottoman and Russian dataset for third and fourth level categories.

2023b) and Falcon (Almazrouei et al., 2023). These models were chosen to maximize efficiency and exploit the knowledge they have from their large pre-training datasets. Llama-2 utilizes a massive dataset collected from publicly available resources, while Falcon uses the RefinedWeb Dataset (Penedo et al., 2023), a cleaned text dataset from CommonCrawl<sup>5</sup>. Additionally, we included multilingual BERT (mBERT) (Devlin et al., 2019) to explore the capabilities of a smaller-scale model in our experiments.

Due to the memory-intensive nature of the pre-trained large-scale transformer models, we opted for "Llama-2-7b"<sup>6</sup> and "Falcon-7b"<sup>7</sup> variants. For smaller scale model we chose "mBERT-base"<sup>8</sup> variant for multilingual BERT. Llama-2-7b and Falcon-7b are the smallest variants available for these model families. We used Nvidia RTX 3090 and 4090 GPUs, both of which have 24 GB of vRAM. For efficient training on our compute resources, we opted for 4-bit quantization (Dettmers et al., 2022). In addition, we did not train all model parameters, but instead used LoRa (Hu et al., 2022) and trained LoRa parameters only. For mBERT, we opted for linear probing and trained only the classification head, keeping the backbone frozen. We also set the floating point type as bfloat16.

In training, we used 2048 context window, a batch size of 4, and conducted a hyperparameter search on learning rate (ranging between  $1.75 \times 10^{-5}$  and  $2 \times 10^{-4}$ ) and weight decay (ranging between  $10^{-4}$  and  $10^{-2}$ ) values.

In addition to LLMs, we use a classical "bag of word naive Bayes" (BoW+NB) model to give a simpler baseline for models prior to LLMs. Each article in the dataset is represented by a binary vector, whose size corresponds to the number of unique words in the dataset of the corresponding classification task. For Russian, the dictionary size varies for different classification tasks, ranging from 163,615 to 242,305 words. For Ottoman, it spans from 55,507 to 179,873 words. To achieve classification tasks, we used a multinomial naive Bayes classifier.

**Fine-tuning over Chunks.** Since articles can exceed 2048 tokens (our max. input length), we

<sup>5</sup><https://commoncrawl.org/>

<sup>6</sup><https://huggingface.co/meta-llama/Llama-2-7b-hf>

<sup>7</sup><https://huggingface.co/tiiuae/falcon-7b>

<sup>8</sup><https://huggingface.co/bert-base-multilingual-cased>

Dataset	Category	Models	Acc.	mAP (macro)	mAP (weighted)	F1 (macro)	F1 (weighted)
Ottoman L1	-	Llama-2-7b	77.99	79.15	<b>81.54</b>	74.54	77.07
		Falcon-7b	77.93	77.93	81.08	<b>82.44</b>	62.65
		mBERT	78.91	78.91	81.00	71.04	<b>77.65</b>
		BoW + NB	<b>79.53</b>	<b>79.53</b>	78.75	64.75	68.05
Russian L1	-	Llama-2-7b	<b>79.11</b>	<b>83.44</b>	<b>87.40</b>	75.26	<b>79.89</b>
		Falcon-7b	77.22	75.10	80.22	58.79	71.79
		mBERT	78.92	80.53	84.30	<b>78.56</b>	74.01
		BoW + NB	77.35	71.45	75.81	70.67	74.85
Ottoman L2	literary text type	Llama-2-7b	53.61	50.05	50.23	51.5	53.04
		Falcon-7b	52.94	50.23	50.23	50.41	50.41
		mBERT	55.89	31.16	65.39	30.84	52.78
		BoW + NB	<b>82.92</b>	<b>87.31</b>	<b>92.50</b>	<b>78.27</b>	<b>83.84</b>
Russian L2	literary text type	Llama-2-7b	68.50	<b>54.52</b>	<b>88.25</b>	<b>39.13</b>	<b>71.25</b>
		Falcon-7b	50.00	37.60	76.96	24.43	53.95
		mBERT	7.41	35.03	73.65	3.45	1.02
		BoW + NB	<b>69.19</b>	38.95	85.95	24.44	65.69

Table 3: Model performances for first-level (L1) and second-level (L2) single-label classification.

trained the model over chunks of articles for those that exceed this length, i.e., we applied chunking with non-overlapping sliding windows of length 2048.

**Chunked Inference.** At test/inference time, we applied chunking to the input article (at the same size as the trained model’s max. input length, 2048), obtained the classification probabilities for each chunk, and then applied average pooling class-wise to obtain the final probabilities of the article.

**Dataset.** For all classification tasks, we used stratified sampling to split the dataset into training (70%), validation (15%), and test (15%) sets. We eliminated the classes with less than 10 samples before splitting and did not use them in training. In addition, we did not conduct experiments on datasets where the sample size of the test is less than 15 instances.

**Metrics.** To evaluate the models, we reported accuracy (percentage of correct predictions), mean average precision (mAP), and F1 scores for single-label tasks. For multi-label tasks, we provided mAP, average precision (AP@0.5), average recall (AR@0.5), and average F1 (AF1@0.5) scores, all globally averaged at the 0.5 threshold. We utilized the scikit-learn package (Pedregosa et al., 2011) for both single-label and multi-label metric computations. For the mAP and F1 metrics, we reported both macro averages (unweighted mean across all classes, i.e., each class is treated equally) and weighted average scores (classes are weighted

by their sample size), although we focused on weighted metrics. While we provide this set of metrics, we monitor the mAP (weighted) metric to assess model quality, as it is common to both single-label and multi-label tasks.

## 5.2 Model Results

Along with the datasets, we provide baseline quantitative results obtained by the chosen models.

The model performances for the single-label classification datasets are given in Table 3. For the L1 datasets, all models perform similarly for both Ottoman and Russian. On the other hand, on the L2 datasets, BoW + NB performs significantly better than other models on Ottoman, and also compared to BoW + NB on Russian L2. This might be due to the fact that literary text types of articles especially for Ottoman can be easily categorized by the certain words they have. However, there is no particular model that outperforms others in all cases. These single-label classification tasks are on higher-level datasets, namely L1 and L2, and are easier tasks compared to the lower-level and multi-label counterparts. Although the model specifications greatly vary (e.g. a billion-parameter model vs. a classical BoW + NB model), the results are comparable among the models for the single-label classification tasks.

The results obtained from the multi-label datasets are given in Table 4. Similar to single-label classification results, simple BoW + NB model performs sufficiently well and surpasses its LLM counterparts on Ottoman L2 and L3 datasets.

Dataset	Category	Models	mAP (macro)	mAP (weighted)	AP@0.5	AR@0.5	AF1@0.5
Ottoman L2	cultural discourse subject	Llama-2-7b	30.62	52.02	62.52	58.77	60.59
		Falcon-7b	30.11	51.65	65.74	58.13	61.70
		mBERT	30.74	53.14	57.24	60.08	58.63
		BoW + NB	<b>42.44</b>	<b>64.43</b>	<b>72.43</b>	<b>70.21</b>	<b>71.30</b>
Ottoman L2	cultural discourse type	Llama-2-7b	18.88	66.09	84.67	70.03	76.82
		Falcon-7b	18.57	63.73	84.76	70.30	76.82
		mBERT	19.70	66.44	84.78	<b>70.48</b>	76.97
		BoW + NB	<b>35.13</b>	<b>69.29</b>	<b>87.23</b>	70.08	<b>77.72</b>
Russian L2	cultural discourse subject	Llama-2-7b	<b>46.94</b>	<b>69.60</b>	<b>69.13</b>	<b>57.14</b>	<b>62.57</b>
		Falcon-7b	26.35	45.06	42.65	29.59	34.94
		mBERT	32.54	60.63	44.12	30.61	36.14
		BoW + NB	26.55	36.04	41.34	34.25	37.46
Russian L2	cultural discourse type	Llama-2-7b	<b>43.96</b>	<b>67.23</b>	48.15	<b>46.43</b>	47.28
		Falcon-7b	24.24	48.13	54.17	<b>46.43</b>	50.00
		mBERT	32.51	59.51	54.17	<b>46.43</b>	50.00
		BoW + NB	36.32	49.93	<b>63.80</b>	45.89	<b>53.38</b>
Ottoman L3	philosophy subject	Llama-2-7b	37.05	40.61	44.75	30.68	36.40
		Falcon-7b	36.28	39.90	43.49	21.82	29.06
		mBERT	36.05	39.54	00.00	00.00	00.00
		BoW + NB	<b>55.20</b>	<b>57.54</b>	<b>68.14</b>	<b>37.55</b>	<b>48.42</b>
Ottoman L3	literary movement	Llama-2-7b	33.02	32.04	36.95	06.85	11.56
		Falcon-7b	25.96	28.68	42.85	01.09	02.14
		mBERT	26.11	29.14	00.00	00.00	00.00
		BoW + NB	<b>40.37</b>	<b>43.09</b>	<b>71.42</b>	<b>16.94</b>	<b>27.39</b>
Ottoman L3	literature subject genre	Llama-2-7b	<b>59.18</b>	65.43	65.60	53.36	58.85
		Falcon-7b	50.85	56.47	68.31	55.19	58.58
		mBERT	48.13	55.30	70.59	58.54	64.00
		BoW + NB	55.38	<b>69.13</b>	<b>80.00</b>	<b>73.33</b>	<b>76.52</b>
Ottoman L3	modernization subject	Llama-2-7b	39.13	42.83	49.18	28.39	34.48
		Falcon-7b	34.67	37.86	47.82	12.50	00.00
		mBERT	27.12	38.42	00.00	00.00	00.00
		BoW + NB	<b>54.45</b>	<b>58.68</b>	<b>51.35</b>	<b>36.53</b>	<b>42.69</b>
Ottoman L3	identity type	Llama-2-7b	65.42	66.48	64.22	77.22	68.42
		Falcon-7b	<b>70.45</b>	70.84	67.28	<b>91.22</b>	<b>74.55</b>
		mBERT	56.12	57.64	<b>68.19</b>	60.00	63.83
		BoW + NB	65.94	<b>76.48</b>	53.33	55.81	54.54
Russian L3	philosophy subject	Llama-2-7b	<b>61.68</b>	<b>69.94</b>	78.38	<b>61.70</b>	<b>69.05</b>
		Falcon-7b	30.73	49.42	<b>74.29</b>	55.32	63.41
		mBERT	47.11	65.29	<b>74.29</b>	55.32	63.41
		BoW + NB	49.64	62.62	60.78	42.95	49.99
Russian L3	literature subject genre	Llama-2-7b	40.08	40.07	30.00	13.64	18.75
		Falcon-7b	38.40	38.77	28.57	9.09	13.80
		mBERT	<b>53.55</b>	54.15	00.00	00.00	00.00
		BoW + NB	52.13	<b>56.51</b>	<b>65.62</b>	<b>70.00</b>	<b>67.74</b>
Russian L3	modernization subject	Llama-2-7b	<b>48.39</b>	<b>53.78</b>	31.25	23.81	27.03
		Falcon-7b	36.01	37.51	18.18	9.52	12.5
		mBERT	38.70	36.61	20.00	19.04	19.51
		BoW + NB	40.85	52.13	<b>80.00</b>	<b>64.51</b>	<b>71.42</b>

Table 4: Model performances for second-level (L2) and third-level (L3) multi-label classifications.

For Russian datasets, Llama consistently performs as the best model in almost all tasks with the only exception being Russian L3 literature-subject-genre where BoW + NB outperformed Llama-2. This is probably because of the pre-training data containing a portion of Russian content. Also for mBERT, we see that the metrics at 0.5 threshold namely, AP, AR, and AF1 are 00.00% for several datasets signaling that the model’s confidence scores for multi-label classification tasks do not even attain the 50% cut-off.

Llama-2 and Falcon are mostly on par for Ottoman Turkish datasets, yet for Russian datasets Llama-2 outperforms Falcon in most cases. Several factors affecting this could be that Llama-2-7b and Falcon-7b were pre-trained on a corpus of 2T with 32K vocabulary size and 1.5T tokens with 65K vocabulary size, respectively.

## 6 Discussion

The experiments show that a simple model like BoW+NB can be comparable to prominent models like Llama-2 and Falcon in a low-resource setting. However, it is important to point out that Llama-2 and Falcon, which have billions of parameters, underwent training with a frozen backbone and 4-bit integer quantization. This approach can substantially hurt the performance of the models.

Based on mAP (weighted), Llama-2 is the best-performing model in both languages in level-1 (L1) tasks, where our dataset has the largest sample size and is well-balanced. For level-2 (L2) and level-3 (L3) tasks, BoW+NB outperforms Llama-2, Falcon, and mBERT in Ottoman classification tasks; yet in Russian, the leading model is still Llama-2. In Russian, the leading model remains Llama-2 in two of the three classification tasks of L3, while BoW+NB is the best-performing model in one of the tasks. Here, it is important to note that, as the sample size gets smaller, we observe the performance difference between BoW+NB and Llama-2 decreases. The results show that, in an imbalanced setting, different metrics elect different models. Due to the nature of our dataset, we can observe this more clearly at lower-level (L2, L3) classification tasks.

At lower-level categories, the classification tasks become more challenging, both in terms of generating a sufficient number of instances and establishing coherent and objective labeling across experts. For example, the subcategories

of the Literature/ Movements labels are open to interpretation, over which experts might disagree. These subcategories, which signify the subject of any critical article on literature, are fluid categories that may overlap in many instances.

In our dataset annotation process detailed in [Section 3.2](#), we opted against the traditional cross-check annotation, and did not calculate inter-annotator agreement (IAA) scores due to the specific nature of our annotation task. Our annotation teams consist of experts proficient in the relevant languages and the historical context of the 18<sup>th</sup> and 19<sup>th</sup> centuries. Unlike tasks such as question answering or sentiment analysis, our categorization process is not inherently subjective but requires specialized expertise. Also, comprehensive guidelines and training were provided to annotators, and regular consensus meetings were held to resolve challenging cases. Consequently, we implemented a hierarchical cross-check system: initial annotations by a team member were reviewed by a peer expert, and then revised by the team leader, whose domain knowledge surpassed that of the other annotators, overwriting the previous annotations. This structure ensured high-quality annotations and rendered traditional IAA metrics less meaningful.

## 7 Conclusion

This paper proposes a collection of carefully curated datasets of some prominent historical non-Western languages, along with a set of benchmarks obtained from contemporary state-of-the-art LLMs and the BoW naive Bayes model. We observed the shortcomings of current leading LLMs in handling downstream tasks with low-resource languages that are not included in their pre-training data. Notably, we also observed a minimal performance difference between BoW and LLMs. This indicates a need for further research on LLMs. We believe that the democratization of contemporary NLP efforts requires the inclusion of low-resource languages. Even if the majority of the population does not use these languages, this line of research can provide significant assistance to historians, linguists, and humanities scholars.

## Acknowledgements

This article received funding from the European Research Council (ERC) project NONWESTLIT under the European Union’s Horizon 2020 research



and innovation program (Grant Agreement No. 950513). Dr. Akbas is supported by the “Young Scientist Awards Program (BAGEP)” of Science Academy, Türkiye. We would like to express our sincere gratitude to the following individuals for their invaluable contributions to this project: Mehtap Öztürk, Zeynep Nur Şimşek, Hazal Bozyer, Jennifer Flaherty, Nikita Drozdov, Ilya Kliger and Kiril Zubkov.

## Limitations

There are three major challenges to this work. First, the data collection process requires experts and significant human effort. All articles must be labeled by human experts. They are required to read each article in detail and categorize it accordingly; however, the variance in the experts’ judgments needs to be handled by cross-checking. In addition, there are currently no digital resources or OCR programs for the Ottoman Turkish language that can perform at an acceptable level without excessive human effort.

Second, in terms of NLP research, the dataset itself poses several challenges. The limited amount of data hinders the fine-tuning process. The articles are very long, requiring their division into chunks. Classification is performed by averaging the predicted probabilities of each 2048-token segment, but most of the segments may not have enough content to classify the whole article into the correct category.

The third limitation lies in the cultural field of the period. Although experts pay considerable attention to capturing a wide representation of formal and thematic properties of the period, which is represented in the taxonomy, it is impossible to have an equal or fair distribution among labels. The cultural and aesthetic world has its sui generis complexities that resist taxonomies and quantitative analysis.

## References

- ABBY FineReader Engine - The most comprehensive OCR SDK for software developers. <https://www.abbyy.com/ocr-sdk/>. Accessed: 2024-05-29.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Heslow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- Rachit Bansal, Himanshu Choudhary, Ravneet Punia, Niko Schenk, Jacob L Dahl, and Émilie Pagé-Perron. 2021. How low is too low? a computational perspective on extremely low-resource languages. *arXiv preprint arXiv:2105.14515*.
- Ethem F Can, Fazli Can, Pinar Duygulu, and Mehmet Kalpakli. 2012. Automatic categorization of ottoman literary texts by poet and time period. In *Computer and Information Sciences II: 26th International Symposium on Computer and Information Sciences*, pages 51–57. Springer.
- Fazli Can, Ethem Can, Pinar Duygulu Sahin, and Mehmet Kalpakli. 2013. Automatic categorization of ottoman poems. *Glottology*, 4(2):40–57.
- Jan Christian Blaise Cruz and Charibeth Cheng. 2020. Establishing baselines for text classification in low-resource languages. *arXiv preprint arXiv:2005.02068*.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. Llm.int8(): 8-bit matrix multiplication for transformers at scale. *arXiv preprint arXiv:2208.07339*.
- 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.
- Jacob et al. Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL HLT 2019*, pages 4171–4186.
- Awet Fesseha, Shengwu Xiong, Eshete Derb Emiru, Moussa Diallo, and Abdelghani Dahou. 2021. Text classification based on convolutional neural networks and word embedding for low-resource languages: Tigrinya. *Information*, 12(2):52.
- Amitha Gopidi and Aniket Alam. 2019. **Computational analysis of the historical changes in poetry and prose**. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 14–22, Florence, Italy. Association for Computational Linguistics.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. The flores evaluation datasets for low-resource machine translation: Nepali-english and sinhala-english. *arXiv preprint arXiv:1902.01382*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. **LoRA: Low-rank adaptation of large language models**. In *International Conference on Learning Representations*.

- Alpay Kabacalı. 2000. *Başlangıcından günümüze Türkiye’de matbaa, basın ve yayın*. Literatür Yayıncılık.
- Anthony Kay. 2007. Tesseract: an open-source optical character recognition engine. *Linux Journal*, 2007(159):2.
- Séamus Lankford, Haithem Afli, and Andy Way. 2023. [adaptmllm: Fine-tuning multilingual language models on low-resource languages with integrated llm playgrounds](#). *Information*, 14(12).
- Chaya Liebeskind and Shmuel Liebeskind. 2020. Deep learning for period classification of historical hebrew texts. *Journal of Data Mining & Digital Humanities*, 2020.
- Deborah A Martinsen. 1997. *Literary journals in imperial Russia*. Cambridge University Press.
- Aaron Mueller, Garrett Nicolai, Arya D McCarthy, Dylan Lewis, Winston Wu, and David Yarowsky. 2020. An analysis of massively multilingual neural machine translation for low-resource languages. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 3710–3718.
- Clemens Neudecker and Lotte Wilms. 2013. Data set for refinement of digitised newspapers.
- Kelechi Ogueji, Yuxin Zhu, and Jimmy Lin. 2021. [Small data? no problem! exploring the viability of pretrained multilingual language models for low-resourced languages](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 116–126, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Aleš Pekárek and Marieke Willems. 2012. The europeana newspapers – a gateway to european newspapers online. In *Progress in Cultural Heritage Preservation*, pages 654–659, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. [The RefinedWeb dataset for Falcon LLM: outperforming curated corpora with web data, and web data only](#). *arXiv preprint arXiv:2306.01116*.
- Alec et al. Radford. 2018. Improving language understanding by generative pre-training. *arXiv preprint arXiv:1805.04946*.
- Yashwanth Reddy Regatte, Rama Rohit Reddy Gangula, and Radhika Mamidi. 2020. Dataset creation and evaluation of aspect based sentiment analysis in telugu, a low resource language. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5017–5024.
- Mariia Ryskina. 2022. *Learning Computational Models of Non-Standard Language*. Ph.D. thesis, Carnegie Mellon University.
- Gözde Gül Şahin and Mark Steedman. 2019. Data augmentation via dependency tree morphing for low-resource languages. *arXiv preprint arXiv:1903.09460*.
- Xin Tang, Shanbo Cheng, Loc Do, Zhiyu Min, Feng Ji, Heng Yu, Ji Zhang, and Haiqin Chen. 2018. Improving multilingual semantic textual similarity with shared sentence encoder for low-resource languages. *arXiv preprint arXiv:1810.08740*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Ashish et al. Vaswani. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- David Vincent. 2000. The rise of mass literacy: Reading and writing in modern europe. (*No Title*).
- W3Techs. 2024. [Usage statistics of content languages for websites](#). [Online; accessed 12-February-2024].
- Dimitrios Zaikis and Ioannis Vlahavas. 2023. From pre-training to meta-learning: A journey in low-resource-language representation learning. *IEEE Access*.

## A Dataset Structure

The dataset structure displaying the hierarchy is given in [Figure 2](#) for the Ottoman Turkish dataset and [Figure 3](#) for the Russian dataset.

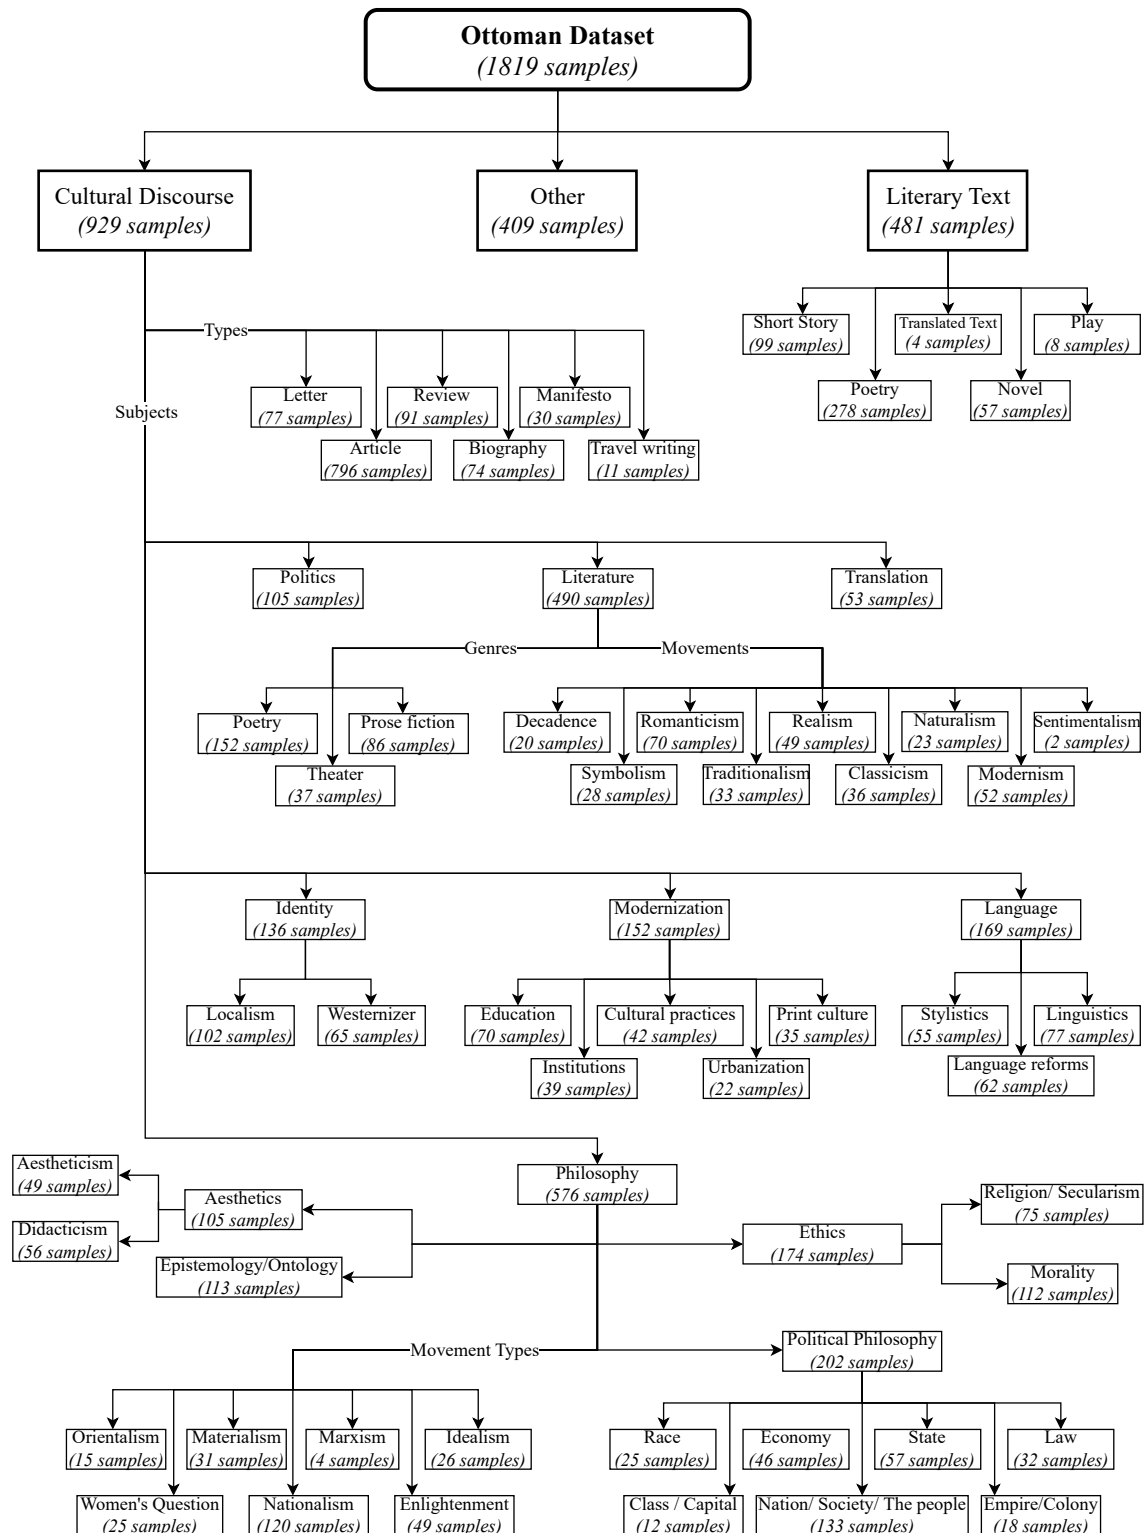


Figure 2: Number of samples in each category for Ottoman Dataset

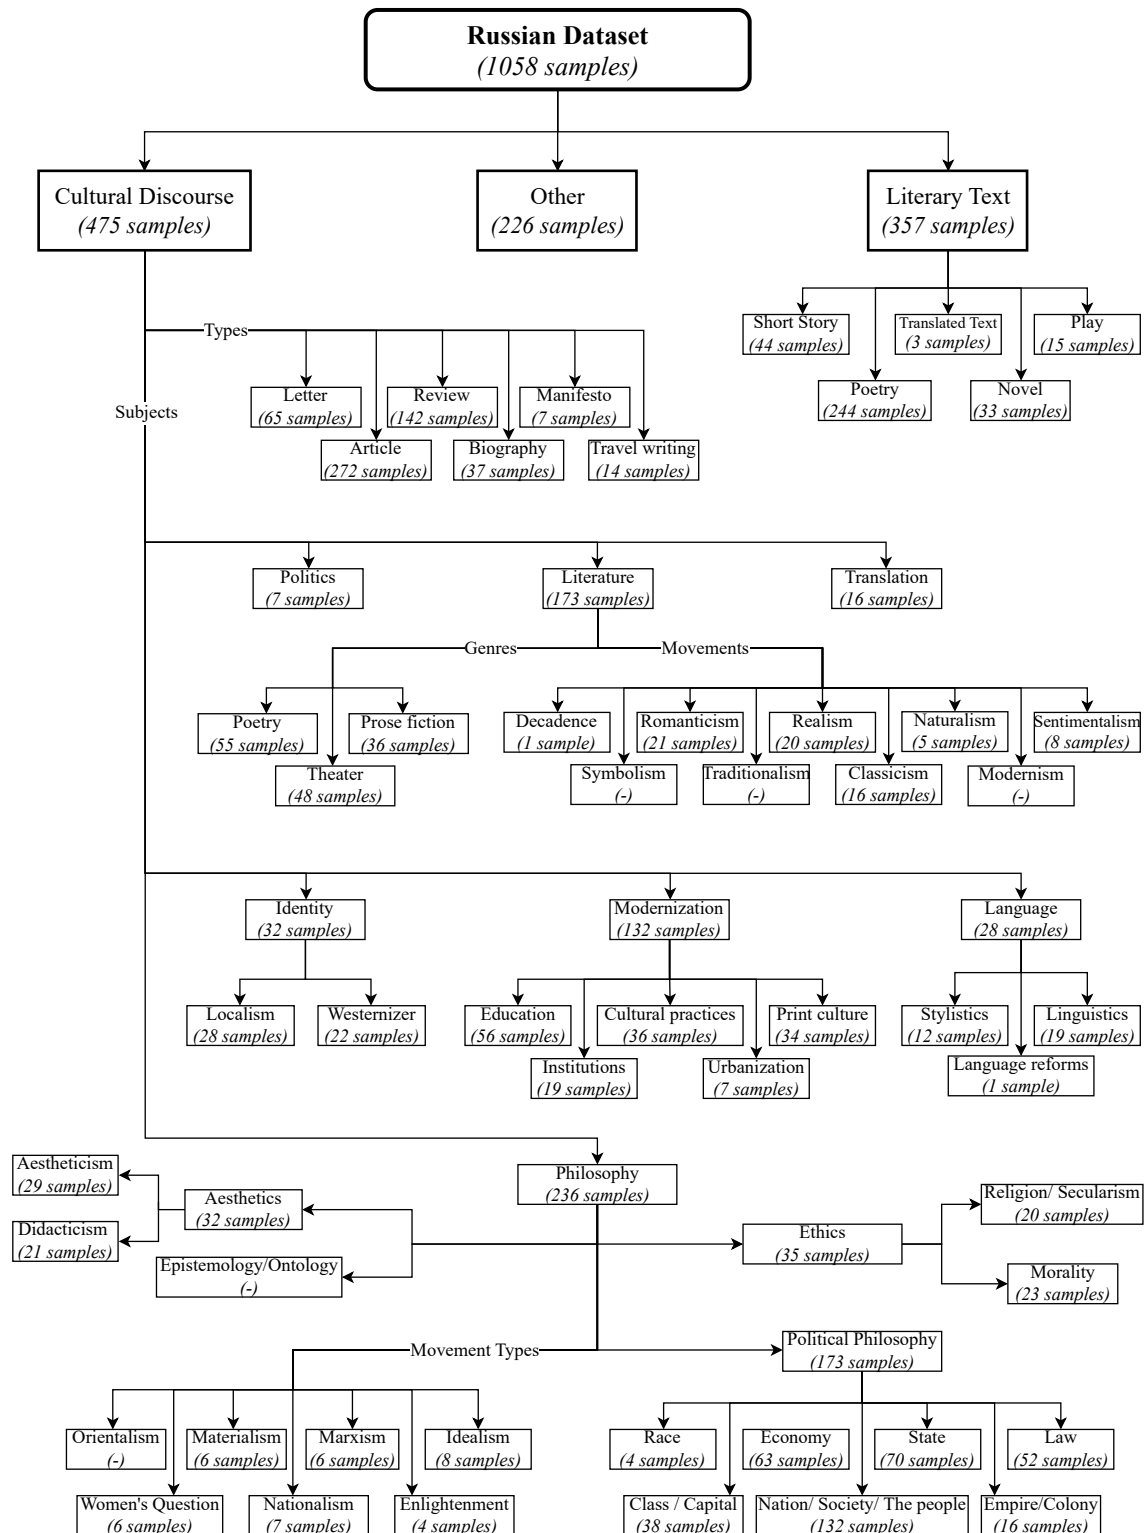


Figure 3: Number of samples in each category for Russian Dataset