

Cross-lingual Named Entity Corpus for Slavic Languages

Jakub Piskorski,^{*} Michał Marcińczuk,[♣] Roman Yangarber[†]

^{*}Polish Academy of Sciences, Poland

[♣]Samurai Labs, Poland

[†]University of Helsinki, Finland

jpiskorski@gmail.com, marcinczuk@gmail.com, roman.yangarber@helsinki.fi

Abstract

This paper presents a corpus manually annotated with named entities for six Slavic languages—Bulgarian, Czech, Polish, Slovenian, Russian, and Ukrainian. This work is the result of a series of shared tasks, conducted in 2017–2023 as a part of the Workshops on Slavic Natural Language Processing. The corpus consists of 5 017 documents on seven topics. The documents are annotated with five classes of named entities. Each entity is described by a category, a lemma, and a unique cross-lingual identifier. We provide two train-tune dataset splits—*single topic out* and *cross topics*. For each split, we set benchmarks using a transformer-based neural network architecture with the pre-trained multilingual models — *XLM-RoBERTa-large* for named entity mention recognition and categorization, and *mT5-large* for named entity lemmatization and linking.

Keywords: named entity, information extraction, natural language processing, lemmatization, Slavic languages

1. Introduction

High-quality recognition and analysis of named entities (NEs) is essential for many information access tasks, such as document retrieval and clustering. It also constitutes a fundamental step in a wide range of natural language processing (NLP) pipelines for higher-level analysis of text, such as information extraction (Grishman, 2019; Huttenen et al., 2002), and anonymization of sensitive data (Templ and Sariyar, 2022).

In this paper, we present a text corpus manually annotated with named entities for six Slavic languages—Bulgarian, Czech, Polish, Slovenian, Russian, and Ukrainian. This corpus results from a series of shared tasks on Named-Entity Recognition, Normalization and Linking (Piskorski et al., 2019, 2021; Yangarber et al., 2023), conducted in 2017–2023 as a part of Workshops on Balto-Slavic Natural Language Processing,¹ whose goal is to stimulate research and foster the creation of tools and resources for these languages, which have over 400 million native speakers.

The corpus consists of 5 017 documents, mainly online news articles, covering seven topics highly debated in the news. The documents are annotated with five categories of named entities: person, organization, location, named mentions of events, and product names—where the last category covers artefacts typically mentioned in the news, e.g., services, awards, cultural artefacts, broadcast media programs, newspapers, legal acts, etc. Each mention of an entity is described with a category, a lemma, a cross-language identifier and positional information. In total, 152

888 named-entity mentions have been annotated, making this the largest cross-lingual NE corpus for Slavic languages.

We also provide two train-tune dataset splits—single-topic-out and cross-topic. For each split, we set benchmarks using a transformer-based neural network architecture with the pre-trained multilingual models — *XLM-RoBERTa-large* for named entity mention recognition and categorization, and *mT5-large* for named entity lemmatization and linking.

The rationale behind the creation and release of the presented corpus is manifold. First, it is a unique new resource in: (a) covering multiple Slavic languages, and documents revolving around the *same topics* across languages, (b) *linking* the names cross-lingually, and (c) providing *base forms* of names. Other existing NER corpora for Slavic languages are strictly monolingual. Secondly, the dataset results from 4 editions of shared tasks on NER for Slavic languages, which attracted a large amount of participants, and is an integration and curation of all datasets used in the tasks with comprehensive annotation guidelines. We have receive inquiries and requests from the research community on the availability of the corpus for training NER, NEL and lemmatisation models that goes beyond the shared task itself. For the first time the integrated and curated corpus is available with positional anchoring—required by many of the approaches applied to the aforementioned shared tasks.

We believe the corpus contributes to fostering research not only on NER, but also name linking and name lemmatization for Slavic languages. Slavic languages exhibit certain challenging phenomena,

¹<https://bsnlp.cs.helsinki.fi/>

like inflection (up to 7 nominal cases), e.g., in newspapers, a significant amount of NE occurrences (i.e., 30-50%) are not nominative forms, and computing base forms of names is a complex task (Piskorski et al., 2009). The corpus covers the domain of online news, which is another factor making it attractive for research on news analysis. The paper is organized as follows. In Section 2 we present related work. In Section 3, we describe the process of creating the corpus, including the taxonomy, annotation process, and corpus statistics. In Section 4, we present the results of evaluation of various benchmark models on the named entity recognition (NER), name lemmatization, and entity linking tasks. We end with conclusions and outlook on future research in Section 6.

2. Related Work

The task of NE recognition has been studied since the 1990's, with a vast amount of research on various approaches, ranging from knowledge-based to machine-learning (Nouvel et al., 2016; Yadav and Bethard, 2018; Jehangir et al., 2023). The research on NER was initially fostered by the NER-related shared tasks in the context of the Message Understanding Conferences (MUCs) (Chinchor, 1998) and the ACE Programme (Dodington et al., 2004). The first *multilingual* NER shared task, which covered several European languages, including Spanish, German, and Dutch, was organized in the context of the the CoNLL conferences (Tjong Kim Sang, 2002; Tjong Kim Sang and De Meulder, 2003). The NE types covered in these campaigns covered mainly the “standard” types: *person*, *organisation*, and *location*. The task of Entity Discovery and Linking (EDL) (Ji et al., 2014, 2015) emerged as a track of the NIST Text Analysis Conferences (TAC). EDL aims to extract entity mentions from a collection of documents in multiple languages (English, Chinese, and Spanish), and to partition the entities into cross-document equivalence classes, by either linking mentions to a knowledge base or directly clustering them.

Related to cross-lingual NE recognition is NE transliteration, i.e., linking NEs across languages that use different alphabets/writing systems. A series of NE Transliteration Shared Tasks were organized as part of NEWS—Named Entity Workshops (Duan et al., 2016), focusing mostly on Indian and Asian languages. In 2010, the NEWS Workshop included a shared task on Transliteration Mining (Kumaran et al., 2010), i.e., mining of names from parallel corpora, in English, Chinese, Tamil, Russian, and Arabic.

Research on NER focusing on Slavic languages includes NER for Croatian (Karan et al., 2013; Ljubešić et al., 2013); NER in Croatian tweets

(Baksa et al., 2017); a manually annotated NE corpus for Croatian (Agić and Ljubešić, 2014); NER in Slovene (Štajner et al., 2013; Ljubešić et al., 2013); a Czech corpus of 11K annotated NEs (Ševčíková et al., 2007); NER for Czech (Konkol and Konopík, 2013); tools and resources for fine-grained annotation of NEs in the National Corpus of Polish (Waszczuk et al., 2010; Savary and Piskorski, 2011); lemmatization of NEs for Polish (Piskorski et al., 2009; Marcińczuk, 2017). Shared tasks on NER for Polish were organized under the umbrella of POLEVAL² (Ogrodniczuk and Kobylński, 2018, 2020) and LESZCZE³ campaigns. Recent shared tasks on NER in Russian include (Starostin et al., 2016; Artemova et al., 2022), with the latter based on NEREL—a Russian dataset for NER and relation extraction, described in Loukachevitch et al., 2021. SemEval 2022 included Task 11: MultiCoNER Multilingual Complex Named Entity Recognition⁴ (Malmasi et al., 2022) and SemEval 2023 included Task 2: MultiCoNER II Multilingual Complex Named Entity Recognition,⁵ (Fetahu et al., 2023) which included Russian and Ukrainian respectively.

A series of Shared Tasks on multilingual NE recognition, normalization and cross-lingual matching for Slavic languages have been organized in the context of the ACL-sponsored workshops on NLP for Slavic languages,⁶ (Piskorski et al., 2017, 2019, 2021; Yangarber et al., 2023) which were the first attempts at such shared tasks covering multiple Slavic languages. The corpus we present in this paper results from merging and curating the corpora from the second (2019), third (2021), and fourth editions (2023) of these shared tasks, where, in contrast to the original datasets, the name mentions are associated with positional information.

3. Named-Entity Annotation

This section describes the end-to-end NE annotation process, including: NE taxonomy, document acquisition, document annotation, the annotation format and the resulting corpus statistics.

3.1. Taxonomy

The Named Entity taxonomy consists of five NE classes: person, organization, location, event and product. The choice of these five entity types has been mainly motivated by the the domain of the

²<http://poleval.pl>

³<https://lepiszzcze.ml/tasks/namedentityrecognition>

⁴https://multiconer.github.io/multiconer_1

⁵<https://multiconer.github.io>

⁶<https://bsnlp.cs.helsinki.fi/>

corpus, namely, online news centred around certain highly-debated topics. In particular, this applies to the inclusion of the less standard entity types, i.e., events (as they are referred to in the news) and products, which cover named mentions of artefacts typically mentioned in the news, e.g., product names, services, awards, cultural artefacts, newspapers, legal acts, etc.

Apart from the entity type, entity annotation also includes information on the base form of the entity, and a cross-lingual ID, so that mentions of the same real-world entity across multiple documents can be assigned to the same unique ID.

More detailed definitions of the five entity categories are provided below.

Person names (PER): Names of real or fictional persons, including initials and pseudonyms, but without including titles, honorifics, and functions or positions. Toponym-based named references to groups of people that have no formal organization unifying them (e.g., “[Ukrainians](#)”) and named mentions of other groups of people that do have a formal organization unifying them (e.g., “[The Royals won](#)”) should be both tagged as PER. Finally, also personal possessives derived from a person’s name should be annotated as PER.

Locations (LOC): All toponyms and geopolitical entities—cities, counties, provinces, countries, regions, bodies of water, land formations, etc.—including named mentions of *facilities*—e.g., stadiums, parks, museums, theaters, hotels, hospitals, transportation hubs, churches, streets, railroads, bridges, and similar facilities. Named mentions of facilities *also* referring to an organization should be tagged with LOC, in “[The Franciszek Raszeia Hospital hired new staff due to the covid pandemics.](#)” the hospital mention should be annotated as LOC.

Organizations (ORG): All organizations, including companies, public institutions, political parties, international, religious, sport, educational and research organizations, etc. Organization designators and mentions of the seat of the organization are considered to be part of the organization name.

Products (PRO): All names of *products and services*, such as electronics, cars, newspapers, web-services, medicines, awards, books, media programmes, initiatives, legal documents, and treaties, e.g., “[Maastricht Agreement](#)”, etc.

Events (EVT): named mentions of incidents, occasions, and events that refer to a specific point of time or a time span, including conferences, concerts, festivals, holidays, e.g., “[Christmas 2024](#)”,

wars, battles, disasters, e.g., “[1943 Gibraltar Liberator AL523 crash](#)”, outbreaks of infectious diseases e.g., “[Covid-19](#)”, and also speculative, and fictive events.

In case of complex *nested* named entities, only the *top-most* entity is annotated.

In case of coordinated phrases mentioning entities, both elements are to be annotated separately, e.g., in “[French and German Ministry of Foreign Affairs](#),” both “[French](#)” and “[German Ministry of Foreign Affairs](#),” are to be annotated.

Detailed Annotation Guidelines are provided in Annex A.

3.2. Document Acquisition

For the creation of the corpus various topics were selected, including, the Covid-19 pandemic, the 2020 USA Presidential elections (USA 2020 Elections), Asia Bibi, which relates to a Pakistani woman involved in a blasphemy case, Brexit, Ryanair, which faced a massive strike, Nord Stream, the controversial Russian-European project, and the Russia-Ukraine war.

For each of the aforementioned topics relevant documents were collected in the following manner. A search query was posed to Google and/or publicly available crawled data repositories, in each of the target languages. The query returned documents in the target language. We removed duplicates, downloaded the HTML—mainly news articles—and converted them into plain text. Since the result of HTML parsing may include not only the main text of a Web page, but also spurious text, some additional manual cleaning was applied when necessary. The resulting set of “cleaned” documents were used to manually select documents for each language and topic for the final datasets. Figure 1 provides the text length histograms for all languages. More than 95% of the texts are not longer than approx. 3K characters.

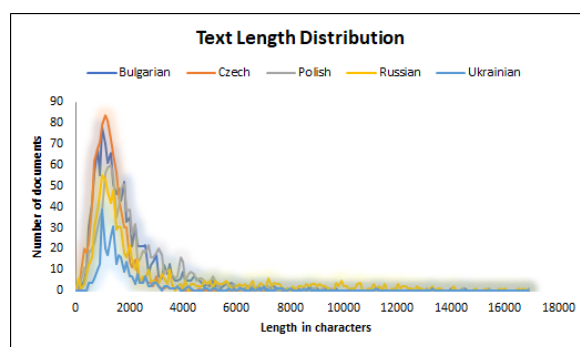


Figure 1: Distribution of the text length for all languages.

3.3. Annotation Platform

Documents were annotated using the Inforex⁷ web-based platform for annotation of text corpora (Marciniuk et al., 2017). The platform allows for sharing a common list of entities, and perform entity-linking semi-automatically: for a given entity, an annotator sees a list of entities of the same type inserted by all annotators and can select an entity ID from the list. Furthermore, it keeps track of all lemmas and IDs inserted for each surface form, and inserts them automatically to speed up the annotation process which in most cases boils down to a confirmation of the proposed values by an annotator. All annotations were made by native speakers. After annotation, we performed *multiple phases* of automatic and manual consistency checks, to reduce annotation errors. The entire process is described in more detail in the next Section.

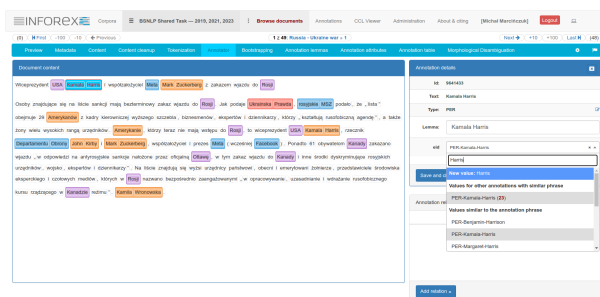


Figure 2: Screenshot of the Inforex Web interface, the tool used for data annotation.

3.4. Annotation Process

The annotation process includes the following steps:

1. all annotators were first provided with detailed annotation guidelines and Q/A sessions were organized,
2. each document was initially annotated by at least one experienced annotator, and a revision was undertaken by master annotator (most experienced annotator for the language) for a given language,
3. to assure high quality annotations, the annotation platform, namely Inforex, was extended with features to automatically suggest the base form and cross-lingual ID for the candidate entities based on all accumulated annotations over time, which turned beneficial for the human annotators,
4. finally, several iterations of data curation (across languages) were carried out with the

annotators to resolve any potential inconsistencies (computed automatically), e.g.:

- same or similar surface forms were assigned different NE types in different documents,
- same surface forms were assigned different base forms,
- same surface form were assigned different cross-lingual IDs.

In the last case, it is worth mentioning that there was an agreed upon and well-defined convention of how cross-lingual IDs are created in order to avoid potential errors.

We encountered some challenges in the annotation process sketched above.

Firstly, there were some disputes regarding what falls under specific NE categories and what does not; this was resolved through various iterations with annotators who annotated texts in different languages, and allowed to converge to a final version of the annotation guidelines.

Secondly, there were disagreements related, in particular, to the ambiguities between ORG vs. PER, and PRO vs. ORG in terms of deciding on the NE type. Here, specific disambiguation rules were introduced in the guidelines, e.g., in cases, in which the local document context and common knowledge does not provide sufficient information to disambiguate the named entity type (e.g., in the phrase *Opel announced that ...* the mention of *Opel* could potentially refer either to an organisation (ORG) or a person (PER), namely, *Adam Opel*, the founder of the company, the more probable interpretation should be considered, i.e., ORG, since Adam Opel died in 1837 and could not announce anything recently, unless the document is a historical one (unlikely). In case both NE type interpretations appear to be equally probable NE type disambiguation rules should be applied. In this particular case, if possible interpretations are ORG or PER, then PER should have priority. The aforementioned NE type ambiguity problems had a direct impact on assigning the correct cross-lingual IDs.

3.5. Annotation Format

For each text document in the dataset a corresponding annotation file exists, which includes for each NE mention a line in the following format:

```
<START> <END> <MENTION> <BASE> <CAT> <ID>
```

where <START>, <END> are the positional information, <MENTION> is the surface form of the mention, <BASE> is the base form of the entity, <CAT> is the the category of the entity (ORG, PER, LOC, PRO,

⁷github.com/CLARIN-PL/Inforex

Input:

[Tusk o Brexicie](#): Nie jest za późno. Nasze serca są otwarte: Nasze serca wciąż są dla was otwarte - powiedział Donald Tusk kilka dni po tym, jak opublikowano sondaż mówiący, że kolejne referendum ws. [Brexitu](#) mogłoby pokazać, że Brytyjczycy wolą zostać w Unii Europejskiej.

Nie jest za późno, aby powstrzymać Brexit, a Wielka Brytania wciąż może zmienić zdanie - powiedział przewodniczący Rady Europejskiej eurodeputowanym w Strasburgu. - Jeśli rząd Wielkiej Brytanii będzie trzymał się swojej decyzji o odejściu, Brexit stanie się rzeczywistością z wszystkimi jej negatywnymi konsekwencjami. Czyż sam David Davis nie powiedział: „jeśli demokracja nie może zmienić zdania, przestanie być demokracją”? - dodał.

Output:

```
0 3 Tusk Tusk PER PER-Donald-Tusk
5 12 Brexicie Brexit EVT EVT-Brexit
90 99 Donald Tusk Donald Tusk PER PER-Donald-Tusk
165 171 Brexitu Brexit EVT EVT-Brexit
189 199 Brytyjczycy Brytyjczycy PER GPE-Great-Britain
211 226 Unii Europejskiej Unia Europejska ORG ORG-EU
257 262 Brexit Brexit EVT EVT-Brexit
265 278 Wielka Brytania Wielka Brytania LOC GPE-UK
326 341 Rady Europejskiej Rada Europejska ORG ORG-EU-Counc
358 367 Strasburgu Strasburg LOC GPE-Strasbourg
379 394 Wielkiej Brytanii Wielka Brytania LOC GPE-UK
434 439 Brexit Brexit EVT EVT-Brexit
511 520 David Davis David Davis PER PER-David-Davis
604 613 Donald Tusk Donald Tusk PER PER-Donald-Tusk
```

Figure 3: An example of a raw document in Polish and a corresponding NE-annotated file.

or EVT), and <ID> is the cross-lingual NE identifier. All the elements are separated by tabs. An example of a raw document in Polish and a corresponding annotation file is shown in Figure 3.

3.6. Statistics

High-level statistics for the entire corpus in terms of the number of documents and text spans annotated are provided in Table 1. The corpus consists of 5017 documents with 152 888 NE mentions in six languages: Polish, Czech, Russian, Bulgarian, Slovene, and Ukrainian. Table 2 provides a comparison of the distribution of NE types across languages. Approximately 80% of the entities are person, organization, or location names. There are, in total, 25 216 unique NE text forms. Of these, 14 831 forms are *hapax legomena* — these forms appear only once in the entire dataset (9.7% of all entities and 58.8% of unique entities). 194 NE forms (approx. 0.8% of all forms) appear at least 100 times in the entire corpus. Figure 4 shows the distribution of NE occurrence for all languages, whereas Table 3 provides the 5 most frequent NE mentions for each language.

The detailed statistics on the distribution of named-entity types, unique surface forms, lemmas, and entity IDs per topic (subcorpus) are provided in Annex B.

4. Baseline models

In this Section we introduce various baseline models and accompanying evaluation results for the

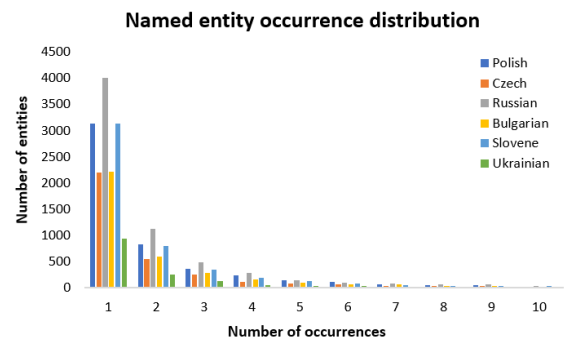


Figure 4: Distribution of named-entity occurrences for all languages.

three tasks at hand: named-entity recognition, name lemmatization, and entity linking.

4.1. Named-Entity Recognition

The baseline model follows the transformer-based neural network architecture presented by Devlin et al. (2019). The model consists of three core components: a pre-trained language model, a dropout layer, and a linear layer that performs token classification. As the pre-trained model, we use the multilingual *XLNet-RoBERTa-large* model (Conneau et al., 2019).⁸ We use the following parameter settings for the training:

Sequence length	256
Dropout	0.2
Epochs	5
Learning rate	5e-6
Optimizer	AdamW
Scheduler	Liner without warmup

We provide baselines for two variants of the dataset split:

- **single topic out** – the test subset is a collection of documents on the topic of US Elections 2020. The train subset contains the remaining six topics (Asia Bibi, Brexit, Nord Stream, Ryanair, Covid, and Russia-Ukraine War).
- **cross topics** – the train and test subsets include documents from each of the seven topics. The test subset contains around 10% of all documents.

The splits intend to benchmark two different aspects of the models. The **single topic out** split evaluates model generality, i.e., how it will perform on new topics. The **cross topics** split evaluates model in-domain performance.

For the **single topic out** split, the baseline models achieved a micro-averaged F-score of 0.8503

⁸<https://huggingface.co/xlm-roberta-large>

topic	PL		CS		RU		BG		SL		UK		Total #doc
	#doc	#an	#doc	#an	#doc	#an	#doc	#an	#doc	#an	#doc	#an	
Brexit	500	16 914	284	5 706	153	4 481	600	16 894	52	2 287	50	1 473	1 639
Asia Bibi	88	1 946	89	1 552	118	2 294	101	2 042	4	89	6	157	406
Nord Stream	151	5 032	161	4 129	150	4 210	130	3 387	74	4 185	40	1 544	706
Ryanair	146	2 386	163	2 502	150	2 306	87	1 176	52	1 558	63	1 194	661
Covid-19	103	2 272	155	2 257	83	4 907	151	3 680	178	5 503	85	2 389	755
US 2020 Elections	66	2 550	85	1 137	170	18 065	151	6 088	143	7 089	83	3 091	698
Russia-Ukraine War	50	1 832	50	1 045	52	1 539	-	-	-	-	-	-	152
Total	1 104	32 932	987	18 328	876	37 802	1 220	33 267	503	20 711	327	9 848	5 017

Table 1: High-level statistics: number of documents (#doc) and text spans annotated (#an) per dataset and language.

Type	PL	CS	RU	BG	SL	UK
PER	7 257	4 430	10 555	8 304	6 831	1 935
LOC	11 603	6 748	12 855	10 838	6 749	3 484
ORG	9 080	4 389	8 960	7 887	4 170	2 787
PRO	2 640	1 356	3 025	1 957	1 251	371
EVT	2 295	1 413	2 337	3 301	1 676	1 271
Total	32 875	18 336	37 732	32 287	20 677	9 848

Table 2: Distribution of named-entity types across languages and statistics on surface forms, lemmas, and unique entities IDs.

across all language categories and categories. The results range from 0.3572 for EVT category to 0.9344 for PER category. The detailed results are presented in Table 4. For all languages, the results range from 0.7933 for RU to 0.9298 for CS (see Table 5).

For the **cross topics** split, the baseline models achieved a micro-averaged F-score of 92.22 across all languages and categories. The results range from 77.07 for PRO category to 96.17 for PER category. The detailed results are presented in Table 6. The results for each language range from 89.13 for RU to 95.89 for PL (see Table 7).

The baseline results for the cross topics split are better than the single topic out split (0.9222 vs. 0.8503 of the F-score) due to a higher overlap of entities between the train and test subsets. Each topic has a set of entities that appears in almost any document on the given topic. For example, the names *Asia Bibi*, *Bibi*, and similar appear solely in the documents about Asia Bibi. In the single-topic out split, the entities specific to US Elections 2020 are less likely to appear in documents on other topics.

Although the presented corpus originates directly from the three editions of the Shared Tasks on Multilingual Named Entity Recognition, Normalization and cross-lingual Matching for Slavic Languages (Piskorski et al., 2019, 2021; Yangarber et al., 2023) a direct comparison of the baseline model results vis-a-vis the results obtained in the aforementioned shared tasks is not straightforward and indicative only since the task has been

defined in a slightly different manner, i.e., it did not focus on the detection of positional information in the text, but on the extraction of unique named entity forms from a given document, where multiple occurrences of the same form are counted only once.

Given that the taxonomy used for annotating the presented corpus is at least to some extent unique, a direct comparison to models trained on other Slavic NER corpora is not straightforward as well. In particular, we do cover the classical PER, ORG and LOC types, however PRO and EVT are somewhat rare vis-a-vis other Slavic NER corpora reported in literature, also in terms of the definitions used.

4.2. Named-Entity Lemmatization

In order to determine the baseline for the lemmatization task, we attempted three different approaches:

- **orth** – the lemma has the same value as the text form. This approach reflects the difficulty of the task, i.e. the amount of named entities that appear in an inflected form.
- **majority** – the lemma is the most frequent lemma in the training corpus. If the text form is not present in the training subset, we take the text form as the lemma.
- **seq2seq** – we trained a sequence-to-sequence neural network using the pre-trained multilingual mT5-large model (Xue et al., 2021).⁹ We use the following parameter settings for the training:

Sequence length	32
Epochs	10
Learning rate	2e-5
Optimizer	Adafactor
Scheduler	Liner without warmup
Weight decay	0.01

⁹<https://huggingface.co/google/mt5-large>

PL		CS		BG		RU		SL		UK	
form	#num	form	#num	form	#num	form	#num	form	#num	form	#num
<i>UE</i>	1 590	<i>EU</i>	702	<i>EC</i>	2 612	<i>США</i>	1 711	<i>ZDA</i>	784	<i>Ryanair</i>	343
<i>Wielkiej Brytanii</i>	895	<i>brexitu</i>	443	<i>Великобритания</i>	2 063	<i>Трампа</i>	754	<i>Trump</i>	725	<i>США</i>	317
<i>Nord Stream 2</i>	790	<i>Ryanair</i>	433	<i>Брекзит</i>	1 920	<i>Ryanair</i>	637	<i>EU</i>	588	<i>ЄС</i>	273
<i>Unii Europejskiej</i>	462	<i>Británie</i>	392	<i>Тръмп</i>	678	<i>Трампа</i>	633	<i>Biden</i>	324	<i>Brexit</i>	235
<i>Brexit</i>	456	<i>Nord Stream 2</i>	257	<i>Лондон</i>	669	<i>Россиа</i>	570	<i>covid-19</i>	187	<i>Трампа</i>	181

Table 3: Top 5 most frequent named-entity forms per language.

Category	Precision	Recall	F ₁	Support
EVT	57.98	25.81	35.72	1 422
LOC	92.28	88.68	90.45	12 247
ORG	73.25	58.24	64.89	6 283
PER	92.09	94.84	93.44	15 936
PRO	69.59	63.64	66.48	2 068
micro avg	87.72	82.51	85.03	37 956
macro avg	86.53	82.51	84.12	37 956

Table 4: Evaluation of the **single topic out** split—summary for all languages.

Language	Precision	Recall	F ₁	Support
PL	90.88	89.53	90.20	2 549
CS	93.98	92.00	92.98	1 137
RU	83.51	75.55	79.33	18 018
BG	93.99	89.52	91.70	6 085
SL	89.77	92.11	90.93	7 082
UK	88.65	77.96	82.96	3 085

Table 5: Evaluation of the **single topic out** split—micro average for each language separately.

For the **cross topics** split, the baseline model achieved an accuracy of 96.13% across all languages. The transformer-based approach outperforms the majority-based lemmatization by less than 4 pp. This is due to the fact that the cross-domain split has a relatively high overlap of entities between the training and testing subsets. The scores for each language are similar and vary between 94 and 97%. The detailed results are presented in Table 8.

For the **single-out topic** split, we can observe a significantly higher difference between the scores obtained by the majority-based approach and the transfer-based model – the difference is more than 16 pp. The machine learning approach achieved an accuracy of 88.89% across all languages. The range of scores for each language is higher and varies from 85 to 93%. The detailed results are presented in Table 9.

4.3. Named-Entity Linking

To establish the baseline for entity linking, we used the same approach as for named entity lemmatization with a sequence-to-sequence neural network. We used exactly the same configuration, except that the output was a cross-language identifier. The average accuracy for the **cross topics** split is 87.84%, and the scores vary from 84 to

Category	Precision	Recall	F ₁	Support
EVT	88.78	90.46	89.61	1 163
LOC	94.94	96.66	95.79	4 577
ORG	86.29	90.69	88.44	3 513
PER	95.63	96.72	96.17	3 752
PRO	75.48	78.72	77.07	954
micro avg	91.03	93.43	92.22	13 959
macro avg	91.10	93.43	92.25	13 959

Table 6: Evaluation of the **cross topics** split—summary for all languages.

Language	Precision	Recall	F ₁	Support
PL	95.30	96.48	95.89	2 986
CS	93.84	96.17	94.99	1 567
RU	87.26	91.09	89.13	3 458
BG	92.20	94.60	93.38	2 872
SL	86.83	89.73	88.25	1 976
UK	92.28	92.36	92.32	1 100

Table 7: Evaluation of the **cross topics** split—micro average, for each language.

91%. The evaluation results are provided in Table 10.

For the **single-out topic** split, the average accuracy for all languages is 68.75%, and the scores vary from 58 to 77%. The evaluation results are provided in Table 11.

It is important to mention that the accuracy figures of the baseline models provided here are not directly comparable with the ones used in the shared tasks (Piskorski et al., 2019, 2021; Yangarber et al., 2023) since a different metric was used, namely, LEA (Moosavi and Strube, 2016).

5. Access

The corpus presented in this paper is publicly available for research purposes at <http://github.com/SlavicNLP/SlavicNER> and <https://bsnlp.cs.helsinki.fi/SlavicNER>. The baseline ML models are publicly available at <https://huggingface.co/SlavicNLP>. For any further use and questions related to the corpus please contact the authors of this paper.

6. Conclusions

This paper describes the construction of a multilingual resource for named entity recognition for Slavic languages. NER is considered the fundamental information extraction (IE) task; it is also an

Language	Orth	Majority	Seq2seq	Support
PL	51.82	93.40	96.68	2 986
CS	56.92	93.80	97.83	1 567
RU	50.30	89.73	94.09	3 458
BG	82.17	95.05	97.45	2 872
SL	57.54	88.30	94.38	1 976
UK	54.09	95.00	98.36	1 100
All	59.25	92.28	96.13	13 959

Table 8: Evaluation of the lemmatization on the **cross topics** split—accuracy for each language.

Language	Orth	Majority	Seq2seq	Support
PL	59.16	77.41	90.13	2 549
CS	54.70	74.67	90.06	1 137
RU	48.79	68.05	86.60	18 018
BG	76.61	82.81	88.51	6 085
SL	52.87	79.20	93.45	7 082
UK	49.23	55.77	91.07	3 085
All	55.24	72.32	88.89	37 956

Table 9: Evaluation of the lemmatization on the **single-out topic** split—accuracy for each language.

essential sub-task in all higher-level IE tasks and in many real-world NLP applications, e.g. (Piskorski and Yangarber, 2013; Linge et al., 2010)

This work results from a series of shared tasks conducted during 2017–2023 as part of the Workshops on Slavic Natural Language Processing. The corpus comprises more than 5K documents on seven topics, and more than 152K entity mentions. We provide a range of detailed statistics, including the distribution of documents, entities, and categories of entities across languages and topics, and the distribution of document lengths across languages.

We provide two splits of the dataset for in-domain (cross-topic split) and cross-domain (single-topic-out) evaluation. To calculate the baseline results for the splits, we trained two models for entity recognition and classification using the pre-trained multilingual *XLM-RoBERTa-large* model. The reported baseline performance is 92.22 and 85.03 micro-averaged F1-score, for the in-domain and cross-domain evaluations, respectively.

The resulting corpus serves as a highly adaptable resource, offering valuable utility for developing and evaluating models designed for a wide range of tasks, including NE recognition, categorization, lemmatization, and the establishment of co-reference connections among entity mentions. This versatility extends to both mono-lingual and cross-lingual applications, making it an invaluable asset for research in the field. The corpus is freely available for research purposes. To our knowledge, this is the first and only multilingual Slavic Named Entity-related corpus with names being cross-lingually linked and provided with base

Language	Seq2seq	Support
PL	89.92	2 986
CS	88.13	1 567
RU	85.13	3 458
BG	89.58	2 872
SL	84.86	1 976
UK	91.09	1 100
All	87.84	13 959

Table 10: Evaluation of the entity linking on the **cross topics** split—accuracy for each language.

Language	Seq2seq	Support
PL	75.13	2 549
CS	77.92	1 137
RU	67.56	18 018
BG	63.60	6 085
SL	76.81	7 082
UK	58.94	3 085
All	68.75	37 956

Table 11: Evaluation of the entity linking on the **single-out topic** split—accuracy for each language.

forms. Other existing NER corpora for Slavic languages are strictly monolingual. Furthermore, for the first time the integrated and curated corpus is available with positional anchoring. We believe the corpus will contribute to fostering research not only on NER, but also name linking and name lemmatization for Slavic languages.

Future extensions of the corpus may include assigning more fine-grained labels to the entities labeled as PRO, since this is a rather broad category. Analogously, the event category could be subdivided into more fine-grained labels, namely, occasions, incidents, natural disasters, phenomena, etc.

7. Acknowledgements

We are grateful to the following people for contributing to the annotation of the data or other activities which contributed to the creation of the presented dataset: Bogdan Babych, Anna Dmitrieva, Zara Kancheva, Olga Kanishcheva, Anisia Katin-skaia, Laska Laskova, Maria Lebedeva, Preslav Nakov, Petya Osenova, Lidia Pivovarova, Senja Pollak, Pavel Pribáň, Ivaylo Radev, Kiril Simov, Marko Robnik-Šikonja, Piotr Rybak, Jan Šnajder, Vasyl Starko, and Josef Steinberger.

Ethics Statement

We find no ethical issues with the current work. We use publicly available resources for all conducted experiments, and release the language resources, which were created in collaboration with colleagues, who were aware of how the data will be used by the research community.

References

- Željko Agić and Nikola Ljubešić. 2014. The SETimes.HR linguistically annotated corpus of Croatian. In *Ninth International Conference on Language Resources and Evaluation (LREC 2014)*, pages 1724–1727, Reykjavík, Iceland.
- Ekaterina Artemova, Maxim Zmeev, Natalia Loukachevitch, Igor Rozhkov, Tatiana Batura, Vladimir Ivanov, and Elena Tutubalina. 2022. [Runne-2022 shared task: Recognizing nested named entities](#). In *Proceedings of the International Conference on Computational Linguistics and Intellectual Technologies “DIALOGUE”*, pages 33–41.
- Krešimir Baksa, Dino Golović, Goran Glavaš, and Jan Šnajder. 2017. Tagging named entities in Croatian tweets. *Slovenščina 2.0: empirical, applied and interdisciplinary research*, 4(1):20–41.
- Nancy Chinchor. 1998. Overview of MUC-7/MET-2. In *Proceedings of Seventh Message Understanding Conference (MUC-7)*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale](#). *CoRR*, abs/1911.02116.
- 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.
- George R. Doddington, Alexis Mitchell, Mark A. Przybocki, Lance A. Ramshaw, Stephanie Strassel, and Ralph M. Weischedel. 2004. [The Automatic Content Extraction \(ACE\) program—tasks, data, and evaluation](#). In *Fourth International Conference on Language Resources and Evaluation (LREC 2004)*, pages 837–840, Lisbon, Portugal.
- Xiangyu Duan, Rafael E. Banchs, Min Zhang, Haizhou Li, and A. Kumaran. 2016. Report of NEWS 2016 machine transliteration shared task. In *Proceedings of The Sixth Named Entities Workshop*, pages 58–72, Berlin, Germany.
- Besnik Fetahu, Sudipta Kar, Zhiyu Chen, Oleg Rokhlenko, and Shervin Malmasi. 2023. [SemEval-2023 task 2: Fine-grained multilingual named entity recognition \(MultiCoNER 2\)](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2247–2265, Toronto, Canada. Association for Computational Linguistics.
- Ralph Grishman. 2019. Twenty-five years of information extraction. *Natural Language Engineering*, 25(6):677–692.
- Silja Huttunen, Roman Yangarber, and Ralph Grishman. 2002. Diversity of scenarios in information extraction. In *Proceedings of the Third International Conference on Language Resources and Evaluation (LREC 2002)*, Las Palmas, Spain.
- Basra Jehangir, Saravanan Radhakrishnan, and Rahul Agarwal. 2023. [A survey on named entity recognition — datasets, tools, and methodologies](#). *Natural Language Processing Journal*, 3:100017.
- Heng Ji, Joel Nothman, and Ben Hachey. 2014. Overview of TAC-KBP2014 entity discovery and linking tasks. In *Proceedings of Text Analysis Conference (TAC2014)*, pages 1333–1339.
- Heng Ji, Joel Nothman, and Ben Hachey. 2015. Overview of TAC-KBP2015 tri-lingual entity discovery and linking. In *Proceedings of Text Analysis Conference (TAC2015)*.
- Mladen Karan, Goran Glavaš, Frane Šarić, Jan Šnajder, Jure Mijić, Artur Šilić, and Bojana Dalbelo Bašić. 2013. CroNER: Recognizing named entities in Croatian using conditional random fields. *Informatica*, 37(2):165.
- Michal Konkol and Miloslav Konopík. 2013. CRF-based Czech named entity recognizer and consolidation of Czech NER research. In *Text, Speech and Dialogue*, volume 8082 of *Lecture Notes in Computer Science*, pages 153–160. Springer Berlin Heidelberg.
- A Kumaran, Mitesh M. Khapra, and Haizhou Li. 2010. Report of NEWS 2010 transliteration mining shared task. In *Proceedings of the 2010 Named Entities Workshop*, pages 21–28, Uppsala, Sweden.
- Jens P Linge, Ralf Steinberger, Flavio Fuart, Stefano Bucci, Jenya Belyaeva, Monica Gemo, Delilah Al-Khudhairy, Roman Yangarber, and Erik van der Goot. 2010. Medisys: medical information system. In *Advanced ICTs for disaster management and threat detection: Collaborative and distributed frameworks*, pages 131–142. IGI Global.

- Nikola Ljubešić, Marija Stupar, Tereza Jurić, and Željko Agić. 2013. Combining available datasets for building named entity recognition models of Croatian and Slovene. *Slovenščina 2.0: empirical, applied and interdisciplinary research*, 1(2):35–57.
- Natalia Loukachevitch, Ekaterina Artemova, Tatiana Batura, Pavel Braslavski, Ilija Denisov, Vladimir Ivanov, Suresh Manandhar, Alexander Pugachev, and Elena Tutubalina. 2021. [NEREL: A Russian dataset with nested named entities, relations and events](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 876–885, Held Online. INCOMA Ltd.
- Shervin Malmasi, Anjie Fang, Besnik Fetahu, Sudipta Kar, and Oleg Rokhlenko. 2022. [SemEval-2022 task 11: Multilingual complex named entity recognition \(MultiCoNER\)](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1412–1437, Seattle, United States. Association for Computational Linguistics.
- Michał Marcińczuk. 2017. Lemmatization of multiword common noun phrases and named entities in Polish. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 483–491, Varna, Bulgaria. INCOMA Ltd.
- Michał Marcińczuk, Marcin Oleksy, and Jan Kočoń. 2017. Inforex - a collaborative system for text corpora annotation and analysis. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017, Varna, Bulgaria, September 2-8, 2017*, pages 473–482. INCOMA Ltd.
- Nafise Sadat Moosavi and Michael Strube. 2016. [Which coreference evaluation metric do you trust? a proposal for a link-based entity aware metric](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 632–642, Berlin, Germany. Association for Computational Linguistics.
- Damien Nouvel, Maud Ehrmann, and Sophie Rosset. 2016. *Named entities for computational linguistics*. John Wiley & Sons.
- Maciej Ogrodniczuk and Łukasz Kobyliński, editors. 2018. [Proceedings of the PolEval 2018 Workshop](#). Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland.
- Maciej Ogrodniczuk and Łukasz Kobyliński, editors. 2020. [Proceedings of the PolEval 2020 Workshop](#). Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland.
- Jakub Piskorski, Bogdan Babych, Zara Kancheva, Olga Kanishcheva, Maria Lebedeva, Michał Marcińczuk, Preslav Nakov, Petya Osenova, Lidia Pivovarova, Senja Pollak, Pavel Přibáň, Ivaylo Radev, Marko Robnik-Sikonja, Vasyl Starko, Josef Steinberger, and Roman Yangarber. 2021. [Slav-NER: the 3rd cross-lingual challenge on recognition, normalization, classification, and linking of named entities across Slavic languages](#). In *Proceedings of the 8th Workshop on Balto-Slavic Natural Language Processing*, pages 122–133, Kiyv, Ukraine. Association for Computational Linguistics.
- Jakub Piskorski, Laska Laskova, Michał Marcińczuk, Lidia Pivovarova, Pavel Přibáň, Josef Steinberger, and Roman Yangarber. 2019. The second cross-lingual challenge on recognition, normalization, classification, and linking of named entities across Slavic languages. In *Proceedings of the 7th Workshop on Balto-Slavic Natural Language Processing*, pages 63–74, Florence, Italy. Association for Computational Linguistics.
- Jakub Piskorski, Lidia Pivovarova, Jan Šnajder, Josef Steinberger, and Roman Yangarber. 2017. The first cross-lingual challenge on recognition, normalization and matching of named entities in Slavic languages. In *Proceedings of the 6th Workshop on Balto-Slavic Natural Language Processing*. Association for Computational Linguistics.
- Jakub Piskorski, Karol Wieloch, and Marcin Sydow. 2009. [On knowledge-poor methods for person name matching and lemmatization for highly inflectional languages](#). *Information retrieval*, 12(3):275–299.
- Jakub Piskorski and Roman Yangarber. 2013. Information extraction: Past, present and future. *Multi-source, multilingual information extraction and summarization*, pages 23–49.
- Agata Savary and Jakub Piskorski. 2011. Language Resources for Named Entity Annotation in the National Corpus of Polish. *Control and Cybernetics*, 40(2):361–391.
- Magda Ševčíková, Zdeněk Žabokrtský, and Oldřich Kruza. 2007. Named entities in Czech: annotating data and developing NE tagger. In *International Conference on Text, Speech and Dialogue*, pages 188–195. Springer.
- Tadej Štajner, Tomaž Erjavec, and Simon Krek. 2013. Razpoznavanje imenskih entitet v

- slovenskem besedilu. *Slovenščina 2.0: empirical, applied and interdisciplinary research*, 1(2):58–81.
- A. S. Starostin, V. V. Bocharov, S. V. Alexeeva, A. A. Bodrova, A. S. Chuchunkov, S. S. Dzhumaev, I. V. Efimenko, D. V. Granovsky, V. F. Khoroshevsky, I. V. Krylova, M. A. Nikolaeva, I. M. Smurov, and S. Y. Toldova. 2016. FactRuEval 2016: Evaluation of named entity recognition and fact extraction systems for Russian. In *Computational Linguistics and Intellectual Technologies. Proceedings of the Annual International Conference “Dialogue”*, pages 688–705.
- Matthias Templ and Murat Sariyar. 2022. [A systematic overview on methods to protect sensitive data provided for various analyses](#). *Int. J. Inf. Secur.*, 21(6):1233–1246.
- Erik Tjong Kim Sang. 2002. [Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition](#). In *Proceedings of the 6th Conference on Natural Language Learning - Volume 20, COLING-02*, pages 1–4, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Erik Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003 - Volume 4, CONLL '03*, pages 142–147, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Jakub Waszczuk, Katarzyna Głowińska, Agata Savary, and Adam Przepiórkowski. 2010. Tools and methodologies for annotating syntax and named entities in the National Corpus of Polish. In *Proceedings of the International Multi-conference on Computer Science and Information Technology (IMCSIT 2010): Computational Linguistics – Applications (CLA'10)*, pages 531–539, Wisła, Poland. PTI.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mt5: A massively multilingual pre-trained text-to-text transformer](#).
- Vikas Yadav and Steven Bethard. 2018. [A survey on recent advances in named entity recognition from deep learning models](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2145–2158, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Roman Yangarber, Jakub Piskorski, Anna Dmitrieva, Michał Marcińczuk, Pavel Přebáň, Piotr Rybak, and Josef Steinberger. 2023. [Slav-NER: the 4th cross-lingual challenge on recognition, normalization, classification, and linking of named entities across Slavic languages](#). In *Proceedings of the 9th Workshop on Slavic Natural Language Processing 2023 (SlavicNLP 2023)*, pages 179–189, Dubrovnik, Croatia. Association for Computational Linguistics.

Appendices

A. Annotation guidelines

The Annotation concerns labeling five types of named entities: persons (PER), organizations (ORG), locations (LOC), events (EVT), and products (PRO).

Each named entity mention annotation should additionally include the lemma of the named entity and an identifier in such a way that detected mentions referring to the same real-world entity should be assigned the same identifier, which we will refer to as cross-lingual ID.

A.1. General Rules

- When assigning the type to a named entity (ORG, LOC, PER, EVT or PRO) in general it is assumed that the local document context and common knowledge is considered and exploited for resolving ambiguities, e.g., in *“Twitter announced revenues for 2018 ...”* the name *“Twitter”* refers to a company (ORG), whereas in the phrase *“I posted it on Twitter”* the name *“Twitter”* refers to a product (PRO), unless explicitly specified otherwise in the remaining part of these guidelines.
- In cases, in which the local document context and common knowledge does not provide sufficient information to disambiguate the named entity type (e.g., in the phrase *“Opel announced that ...”* the mention of *“Opel”* could potentially refer either to an organisation (ORG) or a person (PER), namely, Adam Opel, the founder of the company, the more probable interpretation should be considered, i.e., ORG, since Adam Opel died in 1837 and could not announce anything recently, unless the document is a historical one (unlikely). In case both NE type interpretation appear to be equally probable the following NE type disambiguation rules should be applied. If possible interpretation is either ORG or PER then PER should have priority. If possible interpretation is ORG or PRO then ORG should have priority.

- Sometimes entities are mentioned using a common noun starting with an uppercase letter, e.g., “*Senat*” (Slovene: “Senat”). If it is clear from the context that they refer to a specific entity (e.g. “*Senat Združenih držav Amerike*” – Slovene: “United States Senate”) such mentions are to be treated as named mentions and should be tagged accordingly. Lowercase and lowercase-uppercase mixture mentions of the full name of an entity are to be treated as named mentions, e.g., “*vlada republike slovenije*” (Slovene: “The government of Republic of Slovenia”) and “*vlada Republike Slovenije*” should be considered as mentions of the entity with the full name “*Vlada Republike Slovenije*”.
- Lemmatisation of a named entity mention refers to lemmatisation of the surface form extracted from the input text, e.g., the lemma for “*UE*” and “*Unii Europejskiej*” (Polish: genitive form of “European Union”) is “*UE*” and “*Unia Europejska*” (Polish: nominative form of “European Union”) respectively, whereas the cross-lingual ID assigned to the two aforementioned NE mentions should be the same. Analogously, lemma of a plural mention of an entity is expected to be nominative plural form (e.g., lemma of the word “*Japončích*” (Czech: genitive plural form of “Japanese”) should be “*Japonci*” – Czech: nominative plural form of “Japanese”), whereas lemma of a singular mention should be nominative singular (e.g., lemma of the word “*Japoncem*” (Czech: genitive singular “Japanese”) should be “*Japonec*”).

A.2. Persons (PER)

1. This category covers named references to individual people (e.g., “*Donald Trump*”) or families (e.g. “*Kaczyńscy*”), and certain named references to groups of people.
2. Person names should not include titles, honorifics, and functions/positions. For example, in the text fragment “*CEO Dr. Jan Kowalski*”, only “*Jan Kowalski*” should be annotated as a person name. However, initials and pseudonyms are considered named mentions of person names and should be annotated. Analogously, in the text fragment “*The Prime Minister of the United Kingdom Theresa May*”, only “*Theresa May*” should be tagged as PER, whereas “*United Kingdom*” should be tagged as LOC.
3. Personal possessives derived from a named mention of a person should be annotated and classified as a person. For instance, for

“*Trumpov tweet*” (Croatian - “*Trump’s email*”) it is expected to annotate “*Trumpov*” as PER and have as the base form: “*Trump*”.

4. Toponym-based (e.g., country name adjectives) references to groups of people that are linked to geopolitical entities¹⁰ should also be annotated and tagged as PER, e.g., “*Ukrainians*”. In this context, mentions of a single member belonging to such groups, e.g., “*Ukrainian*” should be assigned the same cross-lingual ID as plural mentions, i.e., “*Ukrainians*”. Furthermore, it should not matter whether “*Ukrainians*” refer to the entire nation, some “unspecified” part thereof or an organisation related to the geopolitical entity. In all these cases PER category should be used and the cross-lingual ID should be the same as assigned to other toponym-based references to the same geopolitical entity, e.g., “*Ukrainians*” and “*Ukraine*” should be assigned the same cross-lingual ID (different mention type, but the same cross-lingual ID).
5. Although continents (“*Europe*”) and geographical regions (“*Eastern Europe*”) are not geo-political entities, as regards named references to groups of people derived from such toponyms the same rule as specified in rule 4 applies by analogy, e.g., “*Europejczycy*” (Polish: “Europeans”) and “*Europa*” (Polish: “Europe”) should be classified as PER and LOC respectively, and both should be assigned the same cross-lingual ID.
6. Named mentions of other groups of people that do have a formal organization unifying them should be tagged as PER and associated with the same cross-lingual ID as the mentions of the corresponding organisation, e.g., in the phrase “*Spartané vyhráli...*” (Czech: “Spartans won...”) the mention of “*Spartané*” should be tagged as PER and have the same cross-lingual ID as the corresponding sport team, e.g., “*AC Sparta Praha*” (football club). Analogously, phrases like “*Europoslanci*” (Czech: “Members of the European Parliament”) should be cross-linked with the mentions of the European Parliament (ORG).

¹⁰Geo-Political Entities are considered to be complex entities consisting of a population, a government or some administrative body, and a physical location. In the context of this shared task geo-political entities comprise countries, provinces, states, counties and cities, and suchlike entities. In the context of this task international bodies and organisations are not considered geopolitical entities.

7. Mentions to groups of people that do not have a formal organization unifying them should not be annotated, e.g., phrases like “*Muslims*”, which refers to a religious group not linked to any particular organisation should not be tagged.
8. Fictive persons and characters (e.g., “*James Bond*”) are considered as persons.

A.3. Locations (LOC)

1. This category includes all toponyms (e.g., cities, counties, provinces, regions, bodies of water, geological formations, etc.) and named mentions of facilities, i.e., functional and primarily man-made structures, such as: stadiums, parks, museums, theaters, hotels, hospitals, transportation hubs (e.g. airports, sea ports, train stations), churches, streets, railroads, highways, bridges, tunnels, parkings, and other similar urban and non-urban facilities. For instance, “*Łazienki Królewskie w Warszawie*” (Polish: “Łazienki Park in Warsaw”) should be tagged as LOC, whereas general references to facilities without a concrete location as “*parki w Warszawie*” (Polish: “parks in Warsaw”) should not be annotated, i.e. only “*Warszawie*” (Polish: locative form of “Warsaw”) should be annotated as LOC in the latter case.
2. Even in case of named mentions of facilities that refer to an organization, the LOC tag should be used. For example, in the text “*The Schipol airport has acquired new electronic gates*” the mention “*The Schipol airport*” should be annotated and classified as LOC.
3. By analogy to rule 2 in A.3, toponyms, in particular, country names (e.g., “*Polska*”) that refer to geopolitical entities (physical location, population of the country, respective government, nation, or a sport team representing a country) should be annotated and classified as LOC disregarding the specific named mention role (“*Poland*” refers to a range of concepts) and assigned the same unique cross-lingual ID. In this context, the relevant toponyms and toponym-derived adjectives (see the rule 4 in Section A.2) referring to the same geo-political entity should be assigned the same cross-lingual ID. In all other contexts, i.e., when a country name (or other toponym) is used to refer to an organisation that has no specific link to the respective geopolitical entity (unless it is accidental), e.g. a band named “*Russia*”, it should be tagged as ORG.
4. When recognising named mentions of facilities potential mentions of the location are con-

sidered to be part of the full mention, e.g., the entire phrase “*St. Stephen Church in Istanbul*” should be annotated and tagged as LOC.

A.4. Organizations (ORG)

1. This category covers all kind of organizations such as: political parties, public institutions, government units, non-governmental organizations, international organizations (e.g., European Union, NATO, united Nations, etc.) military organizations, companies, religious organizations, sport teams and organizations, education and research institutions, music groups, entertainment and media organizations, etc.
2. Organization designators and potential mentions of the seat of the organization are considered to be part of the organization name. For instance, in the text fragment “*Citi Handlowy w Poznaniu*” (Polish: “City Handlowy bank in Poznań”), the full phrase “*Citi Handlowy w Poznaniu*” should be annotated.

A.5. Events (EVT)

1. This category covers named mentions of events, including: (a) occasions such as conferences, e.g. “*24. Konferencje Žárovného Zinkování*”, concerts, festivals, holidays, e.g., “*Święta Bożego Narodzenia*” (Polish: “Christmas”), (b) incidents such as wars, battles, and man-made disasters, e.g., “*Katastrofa Czernobylska*” (Polish: “Chernobyl catastrophe”), and (c) natural disasters and phenomena (including for instance: earthquakes, e.g. “*Great Alaska Earthquake*”, volcano eruptions “*Eyjafjallajökull Eruption*”, outbreaks of infectious diseases, e.g. “*Spanish Flu*”, “*ebola*” etc.).
2. Future, speculative and fictive events, e.g., “*Czexit*” or “*Polexit*” are considered as event mentions too.
3. In case a named mention of the event does also refer to a location, then it should be tagged as LOC. For example, in the phrase “*He died in Waterloo, just before the end of the battle*”, “*Waterloo*” should be tagged as LOC, not as EVT. However, the cross-lingual ID assigned to such a mention, i.e., “*Waterloo*” should be the same as in the case of other mentions to the battle, e.g., “*The Battle of Waterloo*”.
4. When annotating named mentions of events the potential mentions of the location are considered to be part of the full mention, e.g., the entire phrase “*2004 Winter Olympics in Canada*” should be annotated as EVT.

A.6. Products (PRO)

1. This category covers product names, including for instance: electronics (e.g., “*Motorola Moto Z Play*”), cars (e.g. “*Subaru Forester XT*”), vehicles (e.g., “*Fiat Panda*”), weapons (e.g., “*Kalashnikov AK-47*”), web-based services (e.g., “*Twitter*”), medicines (e.g., “*Oxycodone*”, “*remdesivir*”), stock (e.g., “*Google stock*”), awards (e.g., “*Nobel Prize*”), books (e.g., “*Harry Potter and the Sorcerer’s Stone*”), software (e.g., “*MS Office*”), films (e.g., “*Gone with the Wind*”), TV programmes (e.g., “*Wiadomości TVP*”), newspapers (e.g., “*The New York Times*”) and other pieces of art, etc.
2. Names of legal documents, e.g., “*dyrektywy 2001/14/we Parlamentu Europejskiego i Rady*” (Polish: “directives 2001/14/we of the European Parliament and Council”), treaties, e.g., “*Traktat Lizboński*” (Polish: “Treaty of Lisbon”), initiatives/programmes (e.g., “*Horizon 2020*”) are also considered product names
3. When a company name is used to refer to a service, e.g., “*na Instagramie*” (Polish for “on Instagram”), the mention of “*Instagramie*” is considered to refer to a service/product and should be tagged as PRO. However, when a company name refers to a service expressing an opinion of the company, it should be tagged as ORG.

A.7. Complex names

1. In case of complex named entities, consisting of nested named entities, only the top-most (longest) entity should be annotated. For example, in the text fragment “*George Washington University*” one should not annotate “*George Washington*”, but the entire name, namely, “*George Washington University*”.
2. In case of coordinated phrases like for instance “*European and British Parliament*” two names should be annotated (as ORG), i.e., “*European*” and “*British Parliament*”. The respective lemmas would be “*European*” and “*British Parliament*”. Furthermore, the IDs assigned to these entity mentions should refer to “*European Parliament*” and “*British Parliament*” respectively.
3. In rare cases, plural forms might have two annotations—e.g., in the phrase “*a border between Irelands*”—“*Irelands*” should be annotated twice with identical lemmas but different IDs.

B. Detailed Statistics

lang	PL	CS	Asia Bibi			
			RU	BG	SL	UK
PER	963	790	985	73	47	51
LOC	553	447	764	623	29	67
ORG	354	264	494	280	10	38
PRO	62	48	50	73	3	1
EVT	14	3	1	13	0	0
Total	1946	1552	2294	2042	89	157
forms	508	303	407	412	51	87
lemmas	412	248	317	360	41	77
entity IDs	273	160	178	230	31	64

Table 12: Distribution of named-entity types across languages and statistics on surface forms, lemmas, and unique entities IDs for the domain *Asia Bibi*.

lang	PL	CS	Brexit			
			RU	BG	SL	UK
PER	3607	1375	1795	4003	720	357
LOC	5269	1678	906	5151	564	643
ORG	5430	1604	1096	4830	542	203
PRO	906	293	153	633	32	18
EVT	1702	756	519	2277	414	252
Total	16914	5706	4481	16894	2287	1473
forms	2820	1112	782	1212	596	234
lemmas	2133	841	568	1103	411	177
entity IDs	1507	582	268	781	287	127

Table 13: Distribution of named-entity types across languages and statistics on surface forms, lemmas, and unique entities IDs for the domain *Brexit*.

lang	PL	CS	Nord Stream			
			RU	BG	SL	UK
PER	692	681	454	404	814	87
LOC	2107	2174	1885	1574	2110	655
ORG	1066	638	1115	747	834	777
PRO	1150	616	750	654	370	8
EVT	17	19	5	8	57	17
Total	5032	4129	4210	3387	4185	1544
forms	845	770	892	504	902	336
lemmas	634	550	583	448	600	244
entity IDs	441	392	320	305	461	175

Table 14: Distribution of named-entity types across languages and statistics on surface forms, lemmas, and unique entities IDs for the domain *Nord Stream*.

lang	<i>Ryanair</i>					
	PL	CS	RU	BG	SL	UK
PER	163	171	74	194	151	36
LOC	1112	1071	1092	526	584	618
ORG	967	1170	1057	359	662	520
PRO	136	77	83	90	148	20
EVT	8	13	0	7	10	0
Total	2386	2502	2306	1176	1558	1194
forms	514	475	400	323	671	187
lemmas	418	400	332	315	519	137
entity IDs	322	306	251	245	426	108

Table 15: Distribution of named-entity types across languages and statistics on surface forms, lemmas, and unique entities IDs for the domain *Ryanair*.

lang	<i>Covid</i>					
	PL	CS	RU	BG	SL	UK
PER	552	583	754	489	1098	268
LOC	507	579	1184	1265	1583	609
ORG	662	399	910	909	1379	758
PRO	175	194	744	211	373	187
EVT	376	502	1309	806	1058	567
Total	2272	2257	4907	3680	5503	2389
forms	688	941	1436	1111	2191	625
lemmas	557	745	1128	1016	1770	509
entity IDs	404	558	793	764	1393	369

Table 16: Distribution of named-entity types across languages and statistics on surface forms, lemmas, and unique entities IDs for the domain *Covid*.

lang	<i>Russia-Ukraine War</i>		
	PL	CS	RU
PER	399	277	279
LOC	964	459	700
ORG	302	196	436
PRO	99	59	107
EVT	68	54	17
Total	1832	1045	1539
forms	723	498	725
lemmas	563	384	594
entity IDs	410	280	493

Table 18: Distribution of named-entity types across languages and statistics on surface forms, lemmas, and unique entities IDs for the domain *Russia-Ukraine War*.

lang	<i>US Elections 2020</i>					
	PL	CS	RU	BG	SL	UK
PER	881	553	6214	3141	4001	1136
LOC	1091	340	6324	1699	1879	892
ORG	299	118	3852	762	743	491
PRO	112	69	1138	296	325	137
EVT	110	66	486	190	137	435
Total	2550	1137	18065	6088	7089	3091
forms	475	378	3614	1124	1606	541
lemmas	349	279	2650	1019	1129	390
entity IDs	270	201	1700	667	833	270

Table 17: Distribution of named-entity types across languages and statistics on surface forms, lemmas, and unique entities IDs for the domain *US Elections 2020*.