

# The GINCO Training Dataset for Web Genre Identification of Documents Out in the Wild

Taja Kuzman, Peter Rupnik, Nikola Ljubešić

Jožef Stefan Institute

Jamova cesta 39, 1000 Ljubljana

taja.kuzman@ijs.si, peter.rupnik@ijs.si, nikola.ljubestic@ijs.si

## Abstract

This paper presents a new training dataset for automatic genre identification GINCO, which is based on 1,125 crawled Slovenian web documents that consist of 650,000 words. Each document was manually annotated for genre with a new annotation schema that builds upon existing schemata, having primarily clarity of labels and inter-annotator agreement in mind. The dataset consists of various challenges related to web-based data, such as machine translated content, encoding errors, multiple contents presented in one document etc., enabling evaluation of classifiers in realistic conditions. The initial machine learning experiments on the dataset show that (1) pre-Transformer models are drastically less able to model the phenomena, with macro F1 metrics ranging around 0.22, while Transformer-based models achieve scores of around 0.58, and (2) multilingual Transformer models work as well on the task as the monolingual models that were previously proven to be superior to multilingual models on standard NLP tasks.

**Keywords:** automatic genre identification, web genres, genre classification schema, web corpora, Slovenian language

## 1. Introduction

With the arrival of the Web, it has become significantly easier to collect very large corpora that fuel innovation and creation of advanced resources and language technologies. However, contrary to the traditionally collected corpora, web corpora are built in an automated way which limits the control over the contents that constitute the final corpus (Baroni et al., 2009). One of the post hoc evaluation methods to investigate the corpus composition and quality, and to enrich the corpus with important metadata is automatic genre identification (AGI). This method focuses on genres as text categories based on the author’s purpose, the socially recognized function of a document and/or the conventional patterns of form, following the definition by Orlikowski and Yates (1994).

As this research is a part of the MaCoCu<sup>1</sup> project that aims to collect large corpora for under-resourced languages, our main purpose is to provide a classifier that would efficiently identify genres in web corpora for Slovenian and other languages, which would allow an in-depth analysis of the quality and composition of the newly provided corpora.

In addition to this, annotating the data with genre is beneficial in many other areas where language technology can be improved by a more fine-grained document typology, such as in part-of-speech tagging (Giesbrecht and Evert, 2009), zero-shot dependency parsing (Müller-Eberstein et al., 2021), automatic summarization (Stewart and Callan, 2009), and machine translation (Van der Wees et al., 2018). Furthermore, numerous genre annotation studies have been conducted with the aim of improving the Information Retrieval (IR) tools (Stubbe and Ringlstetter, 2007; Zu Eissen

and Stein, 2004; Vidulin et al., 2007; Roussinov et al., 2001; Finn and Kushmerick, 2006; Boese, 2005).

Our contributions in this work are as follows. First, we propose a new genre classification schema, based on the previous schemata but modified with the goal of 1. using labels recognizable to the corpora users and 2. achieving high inter-annotator agreement, based on considering lexico-grammatical characteristics in addition to the purpose and form of the text. Second, we present the Genre Identification Corpus – GINCO 1.0 (Kuzman et al., 2021), a realistic dataset randomly sampled from two Slovenian web corpora that also includes noise, multi-genre documents and other web-specific challenges, enabling evaluation of classifiers in realistic conditions. Finally, we perform machine learning experiments over the new dataset and share some interesting insights with the community.

## 2. Related work

Despite the considerable benefits of automatic genre identification (AGI), no established classification exists (Sharoff, 2010). The genre researchers are not consistent in the use of terminology, and they refer to genres, text types, functional text dimensions or registers in different ways (Sharoff, 2018; Egbert et al., 2015; Laippala et al., 2021; Lee, 2002). Furthermore, there is no consensus on the genre definition. Consequently, most studies use their own genre schema, either hierarchical (Stubbe and Ringlstetter, 2007; Egbert et al., 2015) or not (Asheghi et al., 2016; Sharoff, 2018), and applying single (Santini, 2007; Sharoff, 2010) or multiple label annotation (Vidulin et al., 2007; Laippala et al., 2019). Schemata vary significantly regarding the number of classes as well, which range from seven (Santini, 2007; Sharoff, 2010; Lee and Myaeng, 2002) to more than hundred (Roussinov et al., 2001) and almost 300

<sup>1</sup><https://macocu.eu/>

classes (Crowston et al., 2010). Consequently, the annotated corpora are not comparable, and testing of the similarity between them using cross-classification results in low accuracy (see Sharoff et al. (2010)).

Various approaches used to tackle genre identification revealed the task to be challenging due to fundamental difficulties emanating from the genre notion itself. Firstly, the conventions that characterise a genre category are not fixed or static, and instances of genre vary in their prototypicality (see Santini et al. (2010), Santini (2006)). Secondly, web documents sometimes display features of more than one genre, not fitting in discrete classes (see Sharoff (2021) and Repo et al. (2021)). Example of such hybrid texts is a promotion of a product written in a form of a news article. Thirdly, the purpose of the communication cannot always be discerned (see Williams (2000)).

The schemata that serve as a basis of the most recent genre identification studies are the schema of the Corpus of Online Registers of English (CORE) (Egbert et al., 2015) and the Functional Text Dimensions (FTD) approach (Sharoff, 2018) for English and Russian. The approaches were extended to cross-lingual genre classification experiments to eliminate the need for time-consuming and expensive manual annotation of large corpora in other languages. Bulygin and Sharoff (2018) performed genre classification on an Arabic web corpus using machine-translated English and Russian corpora annotated with FTDs. Using smaller manually annotated Finnish, Swedish, and French corpora, Repo et al. (2021) demonstrated that good levels of cross-lingual transfer from the extensive English CORE corpus to other languages can be achieved through a zero-shot learning setting performed with the Transformer-based pre-trained language models. Furthermore, the research showed that these models can achieve strong performance monolingually on small training data. A subsequent study (Rönnqvist et al., 2021) further improved the zero-shot results by performing zero-shot classification with multilingual pre-trained language models, trained on genre corpora in all four languages.

### 3. Dataset construction

#### 3.1. Corpora

We performed genre annotation on two Slovenian web corpora, crawled in different time periods, the slWaC 2.0 corpus (Erjavec and Ljubešić, 2014) from 2014 and the MaCoCu-sl 1.0 corpus from 2021 (Bañón et al., 2022). The corpora were collected by crawling the Slovenian top-level domain (TLD) `.si`, as well as some generic-domain websites highly interlinked with the national TLD. The two corpora have the same structure, and were compiled and preprocessed using the same machinery, i.e., the SpiderLing<sup>2</sup> crawler (Suchomel and Pomikálek, 2012), the jusText<sup>3</sup> tool for

boilerplate removal (Pomikálek, 2011), and the onion<sup>4</sup> tool for identifying duplicates (Pomikálek, 2011). In these web corpora, we consider two paragraphs to be near-duplicates if their intersection of word 5-grams exceeds the 50% threshold. To circumvent spurious topic-genre correlations, and to represent the distribution of genres on the Slovenian internet as closely as possible, we used a random selection of texts.

Documents not deemed to be suitable for genre annotation were labelled with the following *Not Suitable* categories: *Machine Translation*, *Generated Text*, *Not Slovene*, *Encoding Issues*, *HTML Source Code*, *Boilerplate*, *Too Short/Incoherent*, *Too Long* (longer than 5,000 words), *Non-Textual* (no full sentences, e.g., tables, lists), and *Multiple Texts* (multiple texts that cannot be split). This preprocessing step resulted in the “not suitable” part of the Genre Identification Corpus GINCO, which contains 123 texts. Interestingly, although the two Slovenian corpora were equally represented in the dataset, 89% of the unsuitable texts were from the recently crawled corpus and are mostly *Non-textual*, *Machine Translation* or *Too Short/Incoherent*, which points to the conclusion that the quality of Slovenian web texts has deteriorated since the web crawl in 2014. Additionally, the documents that consisted of multiple texts of different genres, where one text is followed by another, so they can be separated, such as news article, followed by comments, were split accordingly into two or more texts. They were assigned IDs that provide information on their origin and the order of the texts, so that they can be further analysed or merged back. At the end, the final dataset – the “suitable part” of GINCO on which the annotation was performed – consisted of 501 texts from the 2014 corpus and 501 texts from the 2021 corpus, i.e., 1002 texts in total.

While we did not perform any experiments in automating the splitting of documents containing multiple texts, our initial experiments in discriminating between suitable and unsuitable texts showed for the problem to be rather hard, achieving a macro F1 of 0.715, while the random baseline macro F1 is around 0.5. At this point, given the class imbalance (123 unsuitable texts vs. 1002 suitable texts), more unsuitable texts are classified as suitable than unsuitable, which we consider not to be satisfactory. Both tasks – eliminating noise from the web corpora and splitting documents containing texts of different genres – are kept for future work.

#### 3.2. Annotation schema

The construction of the annotation schema was based on the following goals:

1. To reach high coverage with respect to real world corpora – to this end, we avoid using schemata that focus on a small set of specific genres (Zu Eissen and Stein, 2004; Santini, 2006; Ashghi et

<sup>2</sup><http://corpus.tools/wiki/SpiderLing>

<sup>3</sup><http://corpus.tools/wiki/Justext>

<sup>4</sup><http://corpus.tools/wiki/Onion>

al., 2016; Lee and Myaeng, 2002; Boese, 2005). Instead, we propose category groups, based on main communicative purposes, identified in previous research (Egbert et al., 2015; Sharoff, 2010; Santini, 2010; Sharoff, 2018).

2. To consider usability for the corpora users and to provide genre labels, recognizable to users as much as possible – to this aim, we avoid abstract labels, such as *content delivery* (Vidulin et al., 2007), *resources* (Stubbe and Ringlstetter, 2007), *recreation* (Sharoff, 2010), *commpuff* (Sharoff, 2018).
3. To provide categories that the annotators can understand with little training and on which they largely agree – to this end, we avoid schemata with a very fine granularity (Crowston et al., 2010; Roussinov et al., 2001; Williams, 2000; Egbert et al., 2015; Berninger et al., 2008) which usually increases the ambiguity (see Sharoff (2010)).

Considering the prospect of extending the research to cross-lingual transfer from large English genre annotated corpus CORE (Egbert et al., 2015) to smaller languages, proposed by Repo et al. (2021) and Rönqvist et al. (2021), we base our annotation schema on the CORE schema which is hierarchical and consists of 8 higher-level categories and more than 50 subcategories. We focus on the subcategories, for which “the labels are intuitive and correspond to register categories in other corpora” (Laippala, 2019). However, the inter-annotator agreement (see Egbert et al. (2015)) and linguistic description of the categories (see Biber and Egbert (2018)) revealed that there is room for improvement. Firstly, due to high granularity of the schema, the results revealed low inter-annotator agreement, as there was no majority agreement, i.e., agreement between at least three of four annotators, on the main category of 31.06% of documents and on the subcategory of 48.98% documents. For at least 10 subgenres there were no instances with majority agreement. Additionally, Biber and Egbert (2018) found that some subcategories are not well-defined linguistically, and that some are highly similar. Thus, we reduce the granularity by merging some of similar categories, such as *News Report* and *Sport Report*, and by including some less frequent categories into broader ones. Secondly, in an attempt to encompass as many relevant web genres as possible, we include some additional genres, identified in a survey of 200 random documents from the slWaC 2.0 Slovenian web corpus (Erjavec and Ljubešić, 2014), which were considered in previous schemata (Roussinov et al., 2001; Boese, 2005; Laippala et al., 2020; Stubbe and Ringlstetter, 2007; Asheghi et al., 2016; Crowston et al., 2010; Rehm et al., 2008).

Our schema does not group categories in the 8 main categories proposed by CORE (*Narrative*, *Opinion*, *Informational Description/Explanation*, *Interactive Dis-*

*ussion*, *How-to/Instructional*, *Informational Persuasion*, *Lyrical*, *Spoken*), since Biber and Egbert (2018) noted that some main categories encompass subgenres that are situationally, linguistically, and functionally different, that some similar subgenres are spread across different main categories and that some main categories overlap extensively and could be merged. Thus, 23 genre categories, based on recognizable labels, purpose, conventional forms and lexico-grammatical features, are grouped into 7 category groups (see Figure 1), 6 of which are based on purpose (*Objective Informative*, *Subjective Reporting*, *Opinion*, *Promotion*, *Dialogue*, *Literature*), whereas 1 is based solely on the form of the text (*Formatted Text*), consisting of visually very easily recognizable genres that could have various purposes, such as *Frequently Asked Questions*. The category groups were introduced to alleviate the annotation decision process, whereas the annotation and machine learning experiments were performed on the level of categories only. For more details on the categories, see their descriptions in the Appendix 1.

Unlike most schemata which consider the case where all web pages should belong to a predefined taxonomy of genres (Lim et al., 2005; Santini, 2007; Sharoff, 2018), we recognize that there might exist web pages that would not fall into any of the predefined genre labels. Following Asheghi et al. (2016), to deal with such web pages, we introduce the 24th category *Other*, intended for texts which purpose is unknown or not covered by other labels. We suggest that during the annotation process, the annotators keep record of possible additional genres that are annotated as *Other*, and if their presence in the corpora is significant, they can be added to the category set.

### 3.3. Annotation procedure

The genre annotation was conducted individually by two annotators – the author of the annotation schema and a second annotator. Both annotators are PhD students in the field of computational linguistics and have a linguistic background. The second annotator received short training on examples that were deemed to be prototypical, and was provided with a decision-tree survey which the annotators used until they sufficiently familiarized themselves with the task.

In addition to this, the annotators followed detailed annotation guidelines with examples of prototypical texts and descriptions of the purpose, form and common lexico-grammatical linguistic characteristics of the genres<sup>5</sup>. With the latter, we diverge from the approaches of the FTD (Sharoff, 2018) and CORE (Egbert et al., 2015) studies. They avoid considering the linguistic characteristics of genres in the annotation process and base the annotation on “the impression the annotators obtained from reading a text” (Sharoff,

<sup>5</sup>The guidelines for multiple languages are available at <https://tajakuzman.github.io/GINCO-Genre-Annotation-Guidelines/>.

Category Group	Objective Informative	Subjective Reporting	Opinion	Promotion	Dialogue	Literature	Formatted Text	Other
Category	News/Reporting	Opinionated News	Opinion/Argumentation	Promotion	Interview	Script/Drama	FAQ	Other
	Announcement		Review	Promotion of a Product	Forum	Lyrical	List of Summaries /Excerpts	
	Information/Explanation			Promotion of Services	Correspondence	Prose		
	Research Article			Invitation				
	Instruction							
	Recipe							
	Call							
	Legal/Regulation							

Figure 1: The Genre Schema

2018) to allow further analyses of linguistic characteristics of genre categories without circularity. We argue that we cannot disregard the possibility that based on annotators' prior knowledge and experience with genres, their decision may be influenced by linguistic characteristics nevertheless. Additionally, by instructing the annotators to base their decisions on presence of concrete characteristics in the text rather than on their impression of the text, the decisions are less subjective, which improves the inter-annotator agreement and consistency. The key linguistic characteristics were chosen based on a preliminary study of 200 texts from the Slovenian web corpora, and although they were not based on the linguistic analyses of the English CORE corpora by Biber and Egbert (2018), they happened to largely correspond to them.

In contrast to the annotation process of CORE (Egbert et al., 2015) which used single-labelling approach, we opted for multi-labelling approach where texts can be annotated with up to three genre labels. In this setting, the primary label is deemed to be the one that is most prevalent and the one that is mainly used for the automated identification, whereas the secondary and tertiary labels provide additional information on the fuzziness of the text, known as a hybrid.

### 3.3.1. Inter-annotator agreement

The annotation was performed in 21 batches of, on average, 50 texts, and the two annotators discussed uncertain cases in meetings, with frequency of meetings decreasing with progression of the annotation campaign. The inter-annotator agreement was calculated on the annotators' labels that were assigned prior to the discussions. The nominal Krippendorff's alpha (Krippendorff, 2018), calculated for agreement on the complete set of 24 categories on the level of primary labels only, reached 0.71, which is above the acceptable threshold of 0.67, defined by Krippendorff (2018). This shows that the proposed schema and annotation procedure improve the reliability of annotation in comparison with the CORE schema, for which Sharoff (2018) reported nominal Krippendorff's alpha of 0.53 on the complete set of 56 subcategories and 0.66 on the 8 main cate-

gories. Regarding the FTD approach, initial research (Sharoff, 2018) achieved higher agreement, reaching Krippendorff's alpha above 0.76, but subsequent research (Suchomel, 2020) reported significantly lower results, i.e., nominal alpha of 0.497, despite using only 9 of initial 18 FTD categories.

Additionally, we performed a manual analysis of agreement where we considered two cases as partial agreement: agreement between the primary label assigned by one annotator with the secondary label assigned by the other, and the agreement between the two primary labels that are a part of the same category group. The analysis revealed perfect agreement on primary labels in 73%, partial agreement in 19%, and no agreement in 8% of texts. Further analysis of cases with no agreement revealed that in 47% of such cases, no agreement was observed due to difficult categorisation of texts with features of many different genres or none at all. Secondly, in 32% of cases there was no agreement due to instances of new phenomena that were not previously sufficiently covered by the annotation guidelines. This shows that as the authors of web texts demonstrate varying levels of expertise or willingness to conform to the conventions, defined by genre experts, (see Sharoff (2021)) annotation guidelines and the schema based on previous work and small preliminary analyses cannot entirely cover all of the diversity found on the Web. Thus, updating the guidelines (and schema) as the task progresses is recommended.

### 3.4. Dataset encoding and availability

The final dataset, Genre Identification Corpus GINCO 1.0 (Kuzman et al., 2021)<sup>6</sup>, consists of the "suitable" and "not suitable" subset, which are released separately in form of a JSON file. The suitable subset consists of texts, manually labelled with genres, and the non-suitable subset comprises documents, considered to be noise and labelled with *Not Suitable* categories. The corpus contains additional metadata, i.e., URL, domain, year, and attribute `hard`, indicating whether a

<sup>6</sup>The corpus is freely available at <http://hdl.handle.net/11356/1467>.

text was hard to annotate. For each of the two subsets the train:dev:test split is encoded in the dataset in a 60:20:20 manner.

The suitable subset has primary, secondary and tertiary labels encoded on three levels of detail – as 24 labels, 21 labels (on this level, the stratified train:dev:test split was performed) and 12 labels. Smaller sets of labels were produced by merging original categories. In this research, we perform experiments mostly on the primary labels only, to which the information from the secondary labels is added in the experiments in the subsection 4.6. We use the sets of 21 and 12 labels. The latter set is used only in the experiments in the subsection 4.7.

Each text instance is encoded as a sequence of paragraphs. In addition to the text (attribute `text`), each paragraph contains information whether the paragraph is considered a near-duplicate by automatic means (the boolean attribute `duplicate`), and finally, manually added information on whether a near-duplicate is informative for the genre identification (the boolean attribute `keep`). In this research, we exploit only the automated information of near-duplicates, as the information of near-duplicate paragraphs to be kept did not prove to be useful during our preliminary experiments. The size of the subsets is described in the Table 1.

subset	texts	pars	words
suitable	1,002	15,050	478,969
suitable (dedup.)	983	7,088	278,075
not suitable	123	3,402	173,778
both subsets	1,125	18,452	652,747

Table 1: Size of the suitable subset (“suitable” – all paragraphs, “suitable (dedup.)” – texts without near-duplicates), not suitable subset, and the sum of both, i.e., the size of the whole GINCO dataset, in terms of number of texts, paragraphs (“pars”) and words.

## 4. Machine learning experiments

### 4.1. Data split

To prepare the dataset for the machine learning experiments, labels with less than 5 instances, namely *FAQ*, *Script/Drama* and *Lyrical*, were merged to the label *Other*. Thus, the experiments were performed on the set of 21 labels. The dataset was then split into train, dev and test in a 60:20:20 manner. Following Asheghi et al. (2014), we ensured that instances from the same web domain were present in only one split to minimize the effect of topic, website design, and the writing style of specific authors. Stratification by the year of crawling and the `hard` parameter was performed in a manual manner, choosing the stratification by primary label that also ensured reasonable stratification by these two additional variables.

### 4.2. Experimental setup

Dev split was used to optimize hyperparameters for different models. The main focus of the hyperparameter search was the number of training epochs to prevent overfitting and optimize micro and macro F1 scores. In accordance to this, 30 epochs for Transformer models and 200 epochs for fastText models were used. For the Transformer models, the sequence length of 512 tokens was used, and the learning rate was set to  $10^{-5}$ .

The models, trained on the train split, were evaluated on the test split via micro F1 and macro F1 to measure both the instance-level and the label-level performance of a specific setup. In each experiment, at least 5 training runs were performed to assure a reasonable sample for measuring statistical significance of differences in performance of specific setups, which was tested via the Mann-Whitney U rank test.

In the remainder of this section, we present the results of our experiments:

- Section 4.3 – a comparison of different technologies at our disposal
- Section 4.4 – the impact of using full texts of web documents instead of texts with near-duplicate paragraphs removed
- Section 4.5 – investigation of the impact of the training data size
- Section 4.6 – the impact of using secondary labels as additional signal
- Section 4.7 – the impact of downcasting the number of labels from 21 to 12

### 4.3. Choice of technology

classifier	micro F1	macro F1
stratified dummy	0.067	0.061
fastText	0.352 ± 0.038	0.217 ± 0.040
fastText + emb.	0.361 ± 0.007	0.219 ± 0.013
XLM-RoBERTa	0.624 ± 0.015	<b>0.579 ± 0.024</b>
SloBERTa	<b>0.629 ± 0.016</b>	0.575 ± 0.037

Table 2: Comparison of classifiers. Deduplicated datasets were used for training and evaluation. ‘fastText + emb’ denotes fastText with pre-trained Slovenian embeddings.

To assess which technology is the most suitable for the automatic genre identification task, we compared fastText (Joulin et al., 2016) and two base-sized Transformer-based pre-trained language models – the monolingual Slovenian SloBERTa (Ulčar and Robnik-Šikonja, 2021) model and multilingual XLM-RoBERTa (Conneau et al., 2019) model. Additionally, as an illustration of the lower bound, a dummy classifier with stratified guessing strategy was implemented. The results of the experiment, summarized in

Table 2, revealed that fastText performs significantly worse than the Transformer models, and the addition of Slovenian embeddings (Ljubešić and Erjavec, 2018) only marginally increased its performance. The monolingual model SloBERTa and the multilingual model XLM-RoBERTa revealed to be the most suitable for the AGI task, with SloBERTa reaching 0.629 in micro F1 and 0.575 in macro F1, and XLM-RoBERTa 0.624 in micro F1 and 0.579 in macro F1. XLM-RoBERTa was included in the comparison following the findings of Repo et al. (2021) where it outperformed monolingual BERT models in this task. However, in our case it did not perform statistically significantly different from SloBERTa.

Two main conclusions can be drawn from these results. For identifying genre, CNN-like classifiers seem not to be up to the task, while Transformer-based models achieve a drastic improvement of the results. Furthermore, similar to results of previous research, multilingual BERT models seem to be as good for modelling the phenomenon as monolingual BERT models. We currently have two hypotheses why this is the case: 1. model pre-training might not bring a lot to the task, but rather the large Transformer model capacity, and 2. genre might be a more generic linguistic task than standard NLP tasks on which monolingual models tend to outperform multilingual models. We plan to test both of these hypotheses in our future research.

Further experiments were performed with the SloBERTa model. However, future experiments in the cross-lingual genre identification will certainly make use of the equally-performing XLM-RoBERTa model.

#### 4.4. Impact of near-duplicate removal

A standard method for pre-processing web-based data is removal of near-duplicates. While the initial experiment presented in Section 4.3 was performed on documents with near-duplicate paragraphs removed, in these experiments we investigate whether keeping all text for classification is beneficial, especially given that for most of the categories there is not much text available.

The results of these experiments are summarized in Table 3. The main finding is that full texts tend to perform better regarding the macro F1 metric, but worse on the micro F1 metric. This result makes sense as it shows for categories with lower results to improve if more text is available per instance, but the instances from the most populous classes seem to perform worse, resulting in an overall worse per-instance performance. Given the overall minor differences in the two setups, and our greater interest in the per-instance performance, i.e., micro F1, we decided to continue our experiments on the deduplicated dataset.

#### 4.5. Impact of training data size

To understand the effect of the training data size on the performance of the models, we repeated training on

dataset	micro F1	macro F1
full	0.607 ± 0.019	<b>0.596 ± 0.033 *</b>
deduplicated	<b>0.629 ± 0.016 ***</b>	0.575 ± 0.037

Table 3: Effect of near-duplicate paragraph removal. Mann-Whitney U rank test was used for p-value estimation, with the hypothesis that the distribution of one metric, marked with an asterisk, was greater than the distribution of another. Asterisks denote p-value: \*\*\* for  $p < 0.001$ , \*\* for  $p < 0.01$ , \* for  $p < 0.05$ .

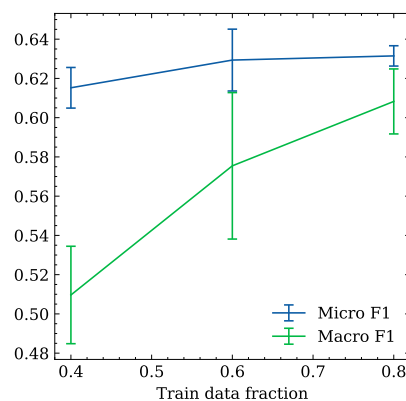


Figure 2: Effect of train data size on micro and macro F1 scores.

train + dev, as well as on a random subset of the training data, comprising of 40% of the whole dataset. In this way, two more datapoints were obtained, with the training sizes being 40%, 60% and 80% of the entire dataset. The effect that this manipulation has on macro and micro F1 scores is shown in Figure 2. The train size correlates positively with both metrics, but the increase is lower after the second train data increase. However, further increasing the training data size should help especially with the low-frequency classes, which is supported by greater improvements obtained on macro F1 than on micro F1.

At this point we analysed the performance of the model on specific categories as well, and compared it with the frequency of each category in the dataset. Figure 3 presents F1 metrics per category, calculated when training on 60% of the data (original train split), ordered by the decreasing frequency of categories. Interestingly, the category frequency does not correlate with the F1 metric. Some of the categories, i.e., *Forum*, *Research Article*, and *Recipe*, are supported by less data than the most frequent categories, but perform better. This indicates that some categories are easier to classify than others. However, one has to bear in mind that the frequency of a category has a direct impact on the size not only of the training data, but also the test data, i.e., the F1 scores of the less represented categories depend more heavily on the (dis)similarity of the few instances in the test split to the few instances in the train split.




Figure 3: Per-category frequency and F1 scores for training on 60% of the deduplicated data. The blue bar depicts the relative frequency of the category, and the purple bar shows the average F1 score from all training runs for a specific category. For reading F1 scores, use bottom  $x$  axis, for relative frequencies, use top  $x$  axis.

#### 4.6. Secondary labels as additional signal

As 188 or 18.7% of the texts are labelled with a secondary category as well, denoting presence of an additional genre, we used this information to inspect whether including the secondary labels in the train split as additional signal would improve the performance of the classifier. The experiments were performed on the original train split (60% of data) of the deduplicated dataset.

A separate training dataset was prepared, where the instances were repeated three times and the last repetition was labelled with the secondary label in an attempt

to augment the performance of the models. With that, we deemed the importance of the primary label to be twice as large as the importance of the secondary label. If there was no secondary label, the instance was repeated three times with the primary label. Given that the instances were repeated three times, the number of training epochs was adapted from 30 to 10 in an attempt to make all experiments comparable. Similarly, we did not adapt the test dataset in any way. As shown in Table 4, the inclusion of secondary labels improved the micro F1 score, while the macro F1 score decreased. The increased micro F1 is not statistically significant, macro F1 decrease, however, is. To conclude, while there might be a positive impact of inclusion of the secondary label, it is minimal and it complicates the setup. Therefore, we opted for using only the primary labels in our final experiment, described in the following subsection.

train labels	micro F1	macro F1
primary	$0.629 \pm 0.016$	$0.575 \pm 0.037$ *
both	$0.635 \pm 0.011$	$0.558 \pm 0.026$

Table 4: Impact of secondary label inclusion in the training dataset. Statistical testing in each column is performed with Mann-Whitney U rank test with p-value encoded with asterisks: \*\*\* for  $p < 0.001$ , \*\* for  $p < 0.01$ , \* for  $p < 0.05$ .

#### 4.7. Number of classes

In the final part of experimentation, we investigate the impact of using a smaller set of labels. To this end, we reduce the number of labels from 21 to 12 labels by merging similar labels (e.g., *Promotion of a Product*, *Promotion of Services*, *Invitation* and *Promotion*) and by adding some less represented labels, such as *Prose*, to the category *Other*.


The use of a smaller label set significantly improved macro and micro F1 scores as shown in Table 5. A positive impact, especially for the categories *Promotion* and *Other*, is also visible in confusion matrices, which are presented in Figure 4a (21 labels) and Figure 4b (12 labels). Here it should be noted once more that the performance of the less represented classes, such as *Legal/Regulation*, *Call*, *Interview* and so on, heavily depends on a very small number of instances in the training and test set. In future work, we plan to double the size of the dataset to provide more training and test examples.

### 5. Conclusion


In this paper we presented a representative dataset of Slovenian web-crawled documents annotated with genre – the freely available Genre Identification Corpus GINCO 1.0 (Kuzman et al., 2021). We introduced a new genre schema which allows annotation with genre of the entire composition of not only web documents,

label set	micro F1	macro F1
21 labels	0.629 ± 0.016	0.575 ± 0.037
12 labels	<b>0.696 ± 0.011 ***</b>	<b>0.668 ± 0.028 ***</b>

Table 5: The effect of using a smaller label set by merging 21 labels to 12 labels. Deduplicated datasets were used for training and evaluation. Statistical testing in each column is performed with Mann-Whitney U rank test with p-value encoded with asterisks: \*\*\* for  $p < 0.001$ , \*\* for  $p < 0.01$ , \* for  $p < 0.05$ .



(a) Original label set (21 labels).



(b) Reduced label set (12 labels).

Figure 4: Confusion matrices for the best performing training run for 21 labels (4a) and 12 labels (4b).

but textual documents in general. Furthermore, we proposed some improvements in the annotation procedure with which we achieved the nominal Krippendorff’s al-

pha (Krippendorff, 2018) of 0.71 which indicates that our approach allows more reliable genre annotation than currently most frequently used approaches.

We performed a series of experiments over the new dataset, revealing that CNN-like classifiers are not up to the task, and the language-specific SloBERTa (Ulčar and Robnik-Šikonja, 2021) Transformer model to be equally potent as the multilingual XLM-RoBERTa (Conneau et al., 2019). Furthermore, we showed that the most reasonable setup for web genre identification is to work with documents with near-duplicate paragraphs removed, using only the dominant, primary labels of texts, and that genre frequency and classification performance do not correlate. When applying this setup to experiments with SloBERTa, we reached 0.629 in micro F1 and 0.575 in macro F1.

While we performed our experiments on a random selection of the Slovenian web, we discarded 10.9% of the data as unsuitable, for which we do not have an efficient classification, or rather elimination approach, and we manually split documents, containing multiple texts with different genres, a task that we did not take on to automate. Once we have these two issues moved out of our way, we can be able to fully claim that we are able to perform high-quality genre identification on a random sample of web data. Nevertheless, this dataset and these experiments still represent the most realistic web-based sample of documents annotated for suitability and genre. Encouraged by positive results, we plan to continue with annotation campaigns to enlarge the Slovenian dataset and to create a Croatian and an English dataset, which will allow cross-lingual experiments.

## Acknowledgements

This work has received funding from the European Union’s Connecting Europe Facility 2014-2020 - CEF Telecom, under Grant Agreement No. INEA/CEF/ICT/A2020/2278341. This communication reflects only the author’s view. The Agency is not responsible for any use that may be made of the information it contains.

This work was also funded by the Slovenian Research Agency within the Slovenian-Flemish bilateral basic research project ”Linguistic landscape of hate speech on social media” (N06-0099 and FWO-G070619N, 2019–2023) and the research programme ”Language resources and technologies for Slovene” (P6-0411).

## 6. Bibliographical References

- Asheghi, N. R., Markert, K., and Sharoff, S. (2014). Semi-supervised graph-based genre classification for web pages. In *Proceedings of TextGraphs-9: The workshop on graph-based methods for natural language processing*, pages 39–47.
- Asheghi, N. R., Sharoff, S., and Markert, K. (2016). Crowdsourcing for web genre annotation. *Language Resources and Evaluation*, 50(3):603–641.



- Baroni, M., Bernardini, S., Ferraresi, A., and Zanchetta, E. (2009). The wacky wide web: a collection of very large linguistically processed web-crawled corpora. *Language resources and evaluation*, 43(3):209–226.
- Berninger, V. F., Kim, Y., and Ross, S. (2008). Building a document genre corpus: a profile of the krys i corpus. In *BCS-IRSG Workshop on Corpus Profiling*, pages 1–10.
- Biber, D. and Egbert, J. (2018). *Register variation online*. Cambridge University Press.
- Boese, E. S. (2005). *Stereotyping the web: genre classification of web documents*. Ph.D. thesis, Citeseer.
- Bulygin, M. and Sharoff, S. (2018). Using machine translation for automatic genre classification in arabic. In *Komp'juternaja Lingvistika i Intellektual'nye Tehnologii*, pages 153–162.
- Crowston, K., Kwaśnik, B., and Rubleske, J. (2010). Problems in the use-centered development of a taxonomy of web genres. In *Genres on the Web*, pages 69–84. Springer.
- Egbert, J., Biber, D., and Davies, M. (2015). Developing a bottom-up, user-based method of web register classification. *Journal of the Association for Information Science and Technology*, 66(9):1817–1831.
- Finn, A. and Kushmerick, N. (2006). Learning to classify documents according to genre. *Journal of the American Society for Information Science and Technology*, 57(11):1506–1518.
- Giesbrecht, E. and Evert, S. (2009). Is part-of-speech tagging a solved task? an evaluation of pos taggers for the german web as corpus. In *Proceedings of the fifth Web as Corpus workshop*, pages 27–35.
- Krippendorff, K. (2018). *Content analysis: An introduction to its methodology*. Sage publications.
- Laippala, V., Kyllönen, R., Egbert, J., Biber, D., and Pyysalo, S. (2019). Toward multilingual identification of online registers. In *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pages 292–297.
- Laippala, V., Rönqvist, S., Hellström, S., Luotolahti, J., Repo, L., Salmela, A., Skantsi, V., and Pyysalo, S. (2020). From web crawl to clean register-annotated corpora. In *Proceedings of the 12th Web as Corpus Workshop*, pages 14–22.
- Laippala, V., Egbert, J., Biber, D., and Kyröläinen, A.-J. (2021). Exploring the role of lexis and grammar for the stable identification of register in an unrestricted corpus of web documents. *Language resources and evaluation*, pages 1–32.
- Laippala, V. (2019). From bits and numbers to explanations—doing research on internet-based big data. In *DATA AND HUMANITIES (RDHUM) 2019 CONFERENCE: DATA, METHODS AND TOOLS*, page 139.
- Lee, Y.-B. and Myaeng, S. H. (2002). Text genre classification with genre-revealing and subject-revealing features. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 145–150.
- Lee, D. (2002). Genres, registers, text types, domains and styles: clarifying the concepts and navigating a path through the bnc jungle. In *Teaching and Learning by Doing Corpus Analysis*, pages 245–292. Brill Rodopi.
- Lim, C. S., Lee, K. J., and Kim, G. C. (2005). Multiple sets of features for automatic genre classification of web documents. *Information processing & management*, 41(5):1263–1276.
- Müller-Eberstein, M., van der Goot, R., and Plank, B. (2021). Genre as weak supervision for cross-lingual dependency parsing. *arXiv preprint arXiv:2109.04733*.
- Orlikowski, W. J. and Yates, J. (1994). Genre repertoire: The structuring of communicative practices in organizations. *Administrative science quarterly*, pages 541–574.
- Rehm, G., Santini, M., Mehler, A., Braslavski, P., Gleim, R., Stubbe, A., Symonenko, S., Tavasani, M., and Vidulin, V. (2008). Towards a reference corpus of web genres for the evaluation of genre identification systems. In *LREC*.
- Repo, L., Skantsi, V., Rönqvist, S., Hellström, S., Oinonen, M., Salmela, A., Biber, D., Egbert, J., Pyysalo, S., and Laippala, V. (2021). Beyond the english web: Zero-shot cross-lingual and lightweight monolingual classification of registers. *arXiv preprint arXiv:2102.07396*.
- Rönqvist, S., Skantsi, V., Oinonen, M., and Laippala, V. (2021). Multilingual and zero-shot is closing in on monolingual web register classification. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 157–165.
- Roussinov, D., Crowston, K., Nilan, M., Kwasnik, B., Cai, J., and Liu, X. (2001). Genre based navigation on the web. In *Proceedings of the 34th annual Hawaii international conference on system sciences*, pages 10–pp. IEEE.
- Santini, M., Mehler, A., and Sharoff, S. (2010). Riding the rough waves of genre on the web. In *Genres on the Web*, pages 3–30. Springer.
- Santini, S. M. (2006). Common criteria for genre classification: Annotation and granularity. In *Workshop on Text-based Information Retrieval (TIR-06), In Conjunction with ECAI 2006, Riva del Garda, 2006*. Citeseer.
- Santini, M. (2007). *Automatic identification of genre in web pages*. Ph.D. thesis, University of Brighton.
- Santini, M. (2010). Cross-testing a genre classification model for the web. In *Genres on the Web*, pages 87–128. Springer.
- Sharoff, S., Wu, Z., and Markert, K. (2010). The web library of babel: evaluating genre collections. In *LREC*. Citeseer.
- Sharoff, S. (2010). In the garden and in the jungle. In *Genres on the Web*, pages 149–166. Springer.

- Sharoff, S. (2018). Functional text dimensions for the annotation of web corpora. *Corpora*, 13(1):65–95.
- Sharoff, S. (2021). Genre annotation for the web: text-external and text-internal perspectives. *Register studies*.
- Stewart, J. G. and Callan, J. (2009). *Genre oriented summarization*. Ph.D. thesis, Carnegie Mellon University, Language Technologies Institute, School of Computer Science.
- Stubbe, A. and Ringlstetter, C. (2007). Recognizing genres. *Proc. Towards a Reference Corpus of Web Genres*.
- Suchomel, V. (2020). Genre annotation of web corpora: Scheme and issues. In *Proceedings of the Future Technologies Conference*, pages 738–754. Springer.
- Van der Wees, M., Bisazza, A., and Monz, C. (2018). Evaluation of machine translation performance across multiple genres and languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Vidulin, V., Luštrek, M., and Gams, M. (2007). Using genres to improve search engines. In *1st International Workshop: Towards Genre-Enabled Search Engines: The Impact of Natural Language Processing*, pages 45–51.
- Williams, Kevin Crowston, M. (2000). Reproduced and emergent genres of communication on the world wide web. *The information society*, 16(3):201–215.
- Zu Eissen, S. M. and Stein, B. (2004). Genre classification of web pages. In *Annual Conference on Artificial Intelligence*, pages 256–269. Springer.
- Ljubešić, Nikola and Erjavec, Tomaž. (2018). *Word embeddings CLARIN.SI-embed.sl 1.0*.
- Pomikálek, Jan. (2011). *Removing boilerplate and duplicate content from web corpora*.
- Suchomel, Vít and Pomikálek, Jan. (2012). *Efficient Web Crawling for Large Text Corpora*.
- Ulčar, Matej and Robnik-Šikonja, Marko. (2021). *Slovenian RoBERTa contextual embeddings model: SloBERTa 2.0*.

## 7. Language Resource References

- Bañón, Marta and Esplà-Gomis, Miquel and Forcada, Mikel L. and García-Romero, Cristian and Kuzman, Taja and Ljubešić, Nikola and van Noord, Rik and Pla Sempere, Leopoldo and Ramírez-Sánchez, Gema and Rupnik, Peter and Suchomel, Vít and Toral, Antonio and van der Werff, Tobias and Zaragoza, Jaume. (2022). *Slovene web corpus MaCoCu-sl 1.0*.
- Alexis Conneau and Kartikay Khandelwal and Naman Goyal and Vishrav Chaudhary and Guillaume Wenzek and Francisco Guzmán and Edouard Grave and Myle Ott and Luke Zettlemoyer and Veselin Stoyanov. (2019). *Unsupervised Cross-lingual Representation Learning at Scale*.
- Erjavec, Tomaž and Ljubešić, Nikola. (2014). *The slWaC 2.0 corpus of the Slovene web*.
- Joulin, Armand and Grave, Edouard and Bojanowski, Piotr and Mikolov, Tomas. (2016). *Bag of tricks for efficient text classification*.
- Kuzman, Taja and Brglez, Mojca and Rupnik, Peter and Ljubešić, Nikola. (2021). *Slovene Web genre identification corpus GINCO 1.0*.

## Appendix 1: Genre Categories

Genre	Description
News/Reporting	an objective text which reports on an event recent at the time of writing or coming in the near future
Announcement	an objective text which notifies the readers about new circumstances, asking them to act accordingly
Instruction	an objective text which instructs the readers on how to do something
Recipe	an objective text which instructs the readers on how to prepare food or drinks
Information/Explanation	an objective text that describes or presents a person, a thing, a concept etc.
Research Article	an objective text which presents research, uses formal language and scientific terms
Call	a text which asks the readers to submit a paper, project proposal, original literary text etc., stating requirements and a deadline
Legal/Regulation	an objective formal text that contains legal terms and is clearly structured
Opinionated News	a subjective text which reports on an event recent at the time of writing or coming in the near future
Opinion/Argumentation	a subjective text in which the authors convey their opinion or narrate their experience. It includes promotion of an ideology and other non-commercial causes.
Review	a subjective text in which authors evaluate a certain entity based on their personal experience
Promotion	a subjective text intended to sell or promote something. This category is used when a promotional text does not fall under Promotion of a Product, Promotion of Services or Invitation.
Promotion of a Product	a subjective text which promotes a product, an application, an accommodation, etc.
Promotion of Services	a subjective text which promotes services of a company
Invitation	a text which invites the readers to participate in an event
Interview	a text consisting of questions posed by the interviewer and answers by the interviewee
Forum	a text in which people discuss a certain topic in form of comments
Correspondence	a text addressed to a person or organization with a form similar to a letter, i.e., including a greeting, a complimentary close etc.
Script/Drama	a literary text that mostly consists of dialogue of characters, stage directions and instructions to the actors
Lyrical	a text that consists of verses
Prose	a literary running text that consists of paragraphs
FAQ	a text in which an author informs the reader through questions and answers
List of Summaries/Excerpts	a text which consists of summaries or excerpts of multiple articles/topics (usually from the article archive page)
Other	a text that has a purpose, not covered by other genre categories, or has no clear purpose

Table 6: Description of genre categories