

FINERWEB: Datasets and Artifacts for Scalable Multilingual Named Entity Recognition

Jonas Golde Patrick Haller Alan Akbik

Humboldt Universität zu Berlin
jonas.max.golde.1@hu-berlin.de

Abstract

Recent multilingual named entity recognition (NER) work has shown that large language models (LLMs) can provide effective synthetic supervision, yet such datasets have mostly appeared as by-products of broader experiments rather than as systematic, reusable resources. We introduce FINERWEB, a dataset-creation pipeline that scales the teacher-student paradigm to 91 languages and 25 scripts. Building on FineWeb-Edu, our approach trains regression models to identify NER-relevant passages and annotates them with multilingual LLMs, resulting in about 225k passages with 235k distinct entity labels. Our experiments show that the regression model achieves more than 84 F1, and that models trained on FINERWEB obtain comparable or improved performance in zero shot transfer settings on English, Thai, and Swahili, despite being trained on 19x less data than strong baselines. In addition, we assess annotation quality using LLM-as-a-judge and observe consistently high scores for both faithfulness (3.99/5) and completeness (4.05/5), indicating reliable and informative annotations. Further, we release the dataset with both English labels and translated label sets in the respective target languages because we observe that the performance of current state-of-the-art models drops by 0.02-0.09 F1 when evaluated using target language labels instead of English ones. We release FINERWEB together with all accompanying artifacts to the research community in order to facilitate more effective student-teacher training for multilingual named entity recognition.

1 Introduction

Named entity recognition (NER) is the task of identifying tokens that belong to a predefined set of classes such as “person” or “location” (Lample et al., 2016; Akbik et al., 2018). In recent years, advances in large language models (LLMs) have enabled zero-shot NER across domains and label

types (Xie et al., 2024; Wang et al., 2025). However, the paradigm of prompting is not an ideal fit for the task of information extraction: it requires issuing separate prompts for each label type, or the generated annotations need postprocessing to align with the original text (Santoso et al., 2024). As a consequence, recent work has leveraged LLMs to generate synthetic training data for smaller, more efficient models such as UniNER (Zhou et al., 2024a), LitSet (Golde et al., 2024), and GLiNER (Zaratiana et al., 2024). These models effectively distill knowledge from LLMs and, despite being orders of magnitude smaller, can outperform their larger teacher models, establishing the state-of-the-art in NER.

However, while autoregressive LLMs such as Qwen3 (Yang et al., 2025), Gemma3 (Team et al., 2025), and GPT-4 (OpenAI et al., 2024) support more than 100 languages, current universal NER models focus mostly on the English language, and no multilingual dataset currently exists to train those models. Existing human-labeled multilingual NER resources also fail to close this gap because they typically provide only one of the two necessary dimensions: (i) broad language coverage or (ii) a rich label set. For instance, datasets such as PAN-X (Pan et al., 2017) cover over 150 languages but restrict the label set to coarse types such as “person”, “location”, and “organization” whereas datasets such as DynamicNER (Luo et al., 2025) support over 100 label types but only support 12 comparatively high-resource languages.

In this paper, we introduce a scalable dataset creation pipeline for multilingual named entity recognition (NER) using large language models (LLMs). The pipeline builds on the FineWeb-Edu methodology (Penedo et al., 2024) and combines automated passage selection and LLM-based annotation to create the dataset. As the outcome of this pipeline, we release FINERWEB, an LLM-annotated dataset covering 91 languages and 25 scripts.

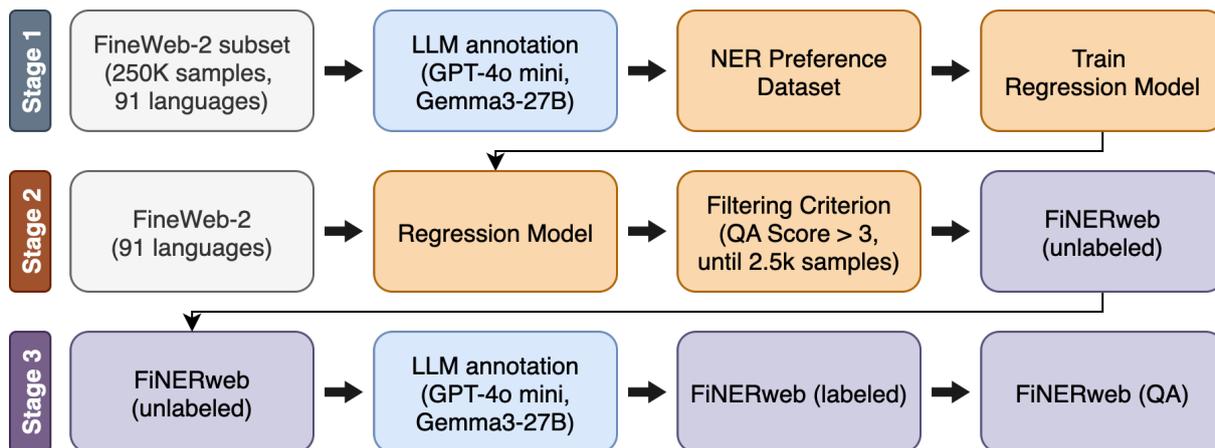


Figure 1: Three-stage pipeline for constructing FiNERWEB: We (1) use a multilingual LLM to generate preference data for identifying high-quality NER examples, (2) train a regression model to filter the FineWeb2 corpus, and (3) annotate the filtered passages with a multilingual LLM.

The pipeline consists of several steps: First, we apply an LLM to rate the usefulness of 1k passages per language for the NER task. These ratings are used to train a multilingual regression model based on XLM-RoBERTa (Conneau et al., 2020). We use the resulting regression model to select 2.5k passages per language from the FineWeb-2 corpus (Penedo et al., 2025). We annotated the selected passages using GPT-4o mini and Gemma3-27B and merge the NER annotations into a single annotation set. At last, we translate all English annotations into the corresponding target languages using the Google Translate API (Google Cloud, 2025).

In our experiments, we demonstrate the usefulness of the resulting dataset by (i) showing that we can train models on selected splits of our dataset and achieve comparable or improved performance compared to current state-of-the-art models and (ii) evaluate the quality of the resulting dataset using LLM-as-a-judge with Qwen3-235B-A22B (Yang et al., 2025). In addition, we analyze the effect of translating the label set into the respective target languages and investigate the long-tail distribution of entity types of the resulting dataset. We summarize our main contributions as follows:

1. We present a scalable pipeline for constructing multilingual NER datasets that combines automated passage selection, LLM-based annotation, and label translation.
2. As the outcome of this pipeline, we release FiNERWEB, an LLM-annotated dataset covering 91 languages and 25 scripts, together with the regression models and their training

data.

3. We conduct extensive experiments to show to usefulness of our dataset including downstream training and LLM-as-a-judge evaluation.
4. We release all artifacts to the research community.¹

2 FiNERweb

In this section, we present the three-stage process used to derive FiNERWEB. Figure 1 illustrates the steps involved. First, we obtain preference training data to train a regression model, which we subsequently apply to filter high-quality passages suitable for NER annotation, e.g., excluding advertisements and other irrelevant content. Finally, we annotate the filtered dataset using two LLMs and merge their outputs to maximize coverage and correctness of entity types.

2.1 Dataset Source

We select FineWeb-2 as our source dataset which covers more than 1,000 languages. However, we cannot simply include all languages, since our work is constrained by the multilingual capabilities of available language models. Therefore, we restrict ourselves to the languages used to train XLM-RoBERTa. This choice ensures compatibility, as the resulting dataset can be directly applied to train or fine-tune XLM-RoBERTa or similar models such as mDeBERTa (He et al., 2021).

¹<https://github.com/whoisjones/fiNERweb>

TASK	ANNOTATION LLM	TRANSFORMER	PRECISION	RECALL	F1
binary	GPT-4o mini	XLM-RoBERTa	<u>0.841</u>	0.842	0.841
		mDeBERTa	0.821	<u>0.827</u>	<u>0.823</u>
	Gemma3 27B	XLM-RoBERTa	0.788	0.761	0.773
		mDeBERTa	0.849	0.536	0.504
multi-class	GPT-4o mini	XLM-RoBERTa	0.667	0.666	0.656
		mDeBERTa	0.582	0.559	0.560
	Gemma3 27B	XLM-RoBERTa	<u>0.637</u>	<u>0.570</u>	<u>0.569</u>
		mDeBERTa	0.595	0.522	0.515

Table 1: Evaluation results (macro-averaged precision, recall and F1 across languages) of our regression models trained on data annotated by different LLMs (4o-mini and Gemma3-27B) and using different transformer architectures. We highlight the best results per setting in bold and the second-best with underlining.

HIGH-QUALITY EXAMPLE (SCORE > 3.5)	LOW-QUALITY EXAMPLE (SCORE < 1.5)
Kraft Foods has taken the Cadbury chocolate brand in a new direction, by combining it with cheese for the first time. The company is bringing together two of its brands and launching Philadelphia with Cadbury, a chilled chocolate spread made from Philadelphia Light and Cadbury chocolate. [...]	Viewing Single Post From: Spoilers for the Week of February 11th. Don't care about Chloe/Taniel/Jen-Jen. Don't care about Sami, really, but hoping that we get some good "SAMANTHA GENE!!" Marlena Death-Stares out of it. STEFANO!! STEFANO, STEFANO, STEFANO!!! [...]

Table 2: English passages selected by our regression model. The high quality passage includes a richer set of entity types, many inferable from context, whereas the low quality passage contains no named entities.

Moreover, we assume that the languages used for XLM-RoBERTa pretraining are also likely to be supported in more recent multilingual LLMs such as Qwen3 or Gemma3.

Further, we use each language in its native type script (for example, Arabic in Arabic script rather than Latin transliteration), and filtered out languages not shared between XLM-R and FineWeb-2. This leaves us with 91 languages in 25 type scripts. As shown in Figure 3, the distribution is dominated by Latin script, followed by smaller shares for scripts such as Cyrillic and Arabic, with a long tail of lower-resource scripts such as Thai and Khmer.

2.2 Stage 1: Selecting High-Quality Passages

In the first step, we aim to filter out noisy and irrelevant texts, such as advertisements or passages that do not contain a diverse range of entity types. This step is necessary because FineWeb-2 is sourced from 96 CommonCrawl snapshots and thus possibly contains many low-quality data points for NER. To do so, we train a regression model that can score the potential quality of a passage for NER annotation.

To construct the training data for the regression

model, we randomly sample 1k examples per language. Each example is obtained by splitting a data point into chunks of 256 tokens using word segmentation models of spaCy (Honnibal et al., 2020), Janome (Uchida, 2015–2025) and Stanza (Qi et al., 2020) from which we sample exactly one chunk to increase diversity. We then use GPT-4o mini and Gemma3-27B to rate the usefulness of each passage for NER on a scale from 1 to 4, following the prompt shown in Figure 8. The scale of scores is defined as:

- **1 point:** None, few or ambiguous entities with little context for reliable identification.
- **2 points:** Clean and coherent text with identifiable entities and minimal noise.
- **3 points:** Diverse entities across domains with sufficient contextual clues.
- **4 points:** Rich, well-contextualized, and noise-free text suitable for NER training and evaluation.

We show example ratings in Table 2 for English and in Appendix C for additional languages.

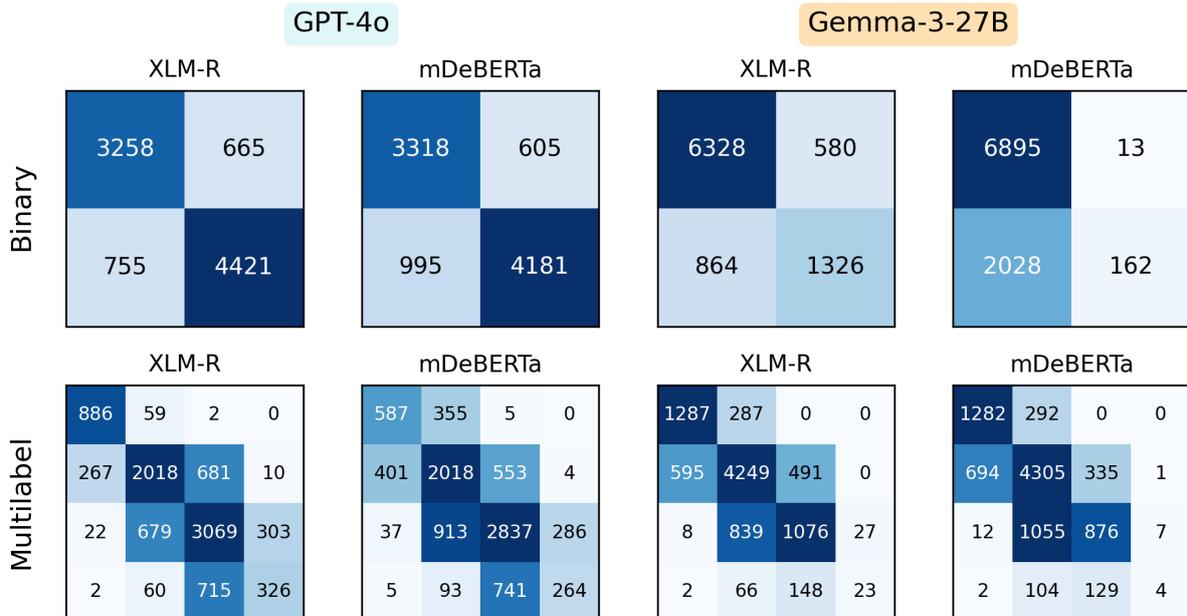


Figure 2: Confusion matrices for regression models reported in Table 1 showing the distribution of prediction errors. GPT-4o mini annotations yield more balanced predictions along the diagonal, whereas Gemma3 annotations rarely give the highest score. Most errors occur close to the diagonal, indicating models avoid severe misclassifications.

Regression Model Training. Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ denote the annotated preference dataset, where x_i is a passage and $y_i \in \{1, 2, 3, 4\}$ is the usefulness score assigned by the LLM. We employ a simple regression model θ to encode x_i and obtain a predicted score $\hat{y}_i \in \mathbb{R}$, which we optimize by minimizing the mean squared error (MSE) between predictions and labels:

$$\mathcal{L}_{\text{reg}} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2.$$

Experimental Setup. We employ two LLMs to annotate the random 1k passages per language, namely GPT-4o mini and Gemma3-27B, and experiment with two different multilingual transformer architectures: XLM-RoBERTa and mDeBERTa. For both models, we use a learning rate of $5e^{-5}$ and a batch size of 64, training for up to 20 epochs with early stopping (patience of 5) in half precision. We allocate 10% of the data as a validation split, using a stratified sample across languages.

We further train two variants: (i) a binary model, in which the models learn to distinguish between useful and non-useful passages (score ≥ 3), and (ii) a multi-class variant, in which the model is trained on all four categories to enable more fine-grained filtering. For each setting, we train three models using different seeds and report the macro-

averaged results over languages across these runs.

Results. We present the classification results of regression model in Table 1 by mapping the continuous predictions \hat{y}_i back to discrete labels by rounding to the nearest integer. We observe that XLM-RoBERTa consistently outperforms mDeBERTa, achieving for instance +2.2 F1 in the binary setting and +8.7 F1 in the multi-class setting when using GPT-4o mini annotations. A similar trend holds for annotations provided by Gemma3-27B. Moreover, models trained on GPT-4o mini annotations consistently yield higher performance compared to those trained on Gemma3 annotations, e.g., 84.1 F1 (binary / GPT-4o-mini annotations / XLM-RoBERTa transformer) versus 77.3 F1 (binary / Gemma3 annotations / XLM-RoBERTa transformer).

We further analyze the model errors in Figure 2. GPT-4o mini annotations are balanced across the diagonal of the confusion matrix, whereas Gemma3-27b annotations are skewed, rarely assigning the maximum score of 4. Although not perfect, most regression model errors are concentrated near the diagonal, suggesting that the models avoid severe misclassifications. In contrast, Gemma3 annotations classify only a small number of passages as highly useful. Based on these findings, we select the binary XLM-RoBERTa regression model trained on GPT-4o mini annotations for Stage 2 of our process.

	NUNER	PILENER	EURO- GLiNER-X	FiNERWEB		
				4o-mini	Gemma	Merged
# Samples	968.4k	45.9k	85.2k	226.0k	225.6k	226.0k
Avg. Text Length	147.8	1063.7	1501.9	1310.9	1308.1	1310.9
Avg. Types/Sample	4.5	20.5	8.4	19.4	21.5	25.4
Avg. Unique Types/Sample	3.3	5.2	3.7	6.6	9.5	11.9
Unique Entity Types	192k	12.6k	20.0k	28.4k	134k	235k
# Languages	1	1	14	91	91	91

Table 3: Comparison of FiNERWEB with universal NER datasets NuNER, PileNER, and Euro-GLiNER-X. We highlight the largest quantities in bold.

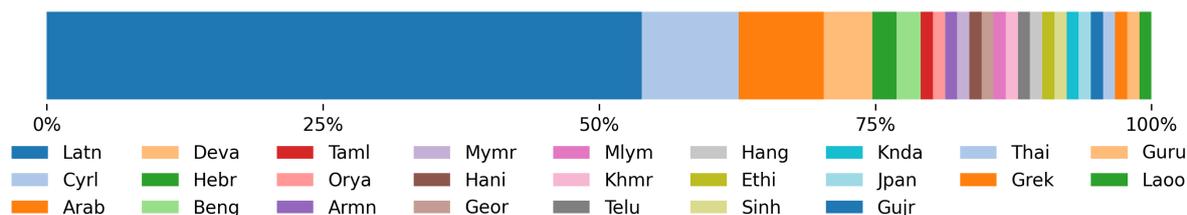


Figure 3: Distribution of scripts in FiNERWEB: Around 50% of the languages use the Latin script, while the dataset covers a long tail of 25 different scripts in total.

2.3 Stage 2: Filtering FineWeb-2

With the regression model from Stage 1, we filter FineWeb-2 for high-quality examples in each target language. We apply identical pre-processing and split each example into chunks of 256 tokens. We then compute the usefulness score using our regression model. If the predicted score is > 0.5 (which is equivalent to the passage scoring ≥ 3). Once we found a useful chunk, we apply an additional quality-control step using a fine-tuned language identification model² to detect whether a passage is not in English as we found some websites to contain embedded English advertisement texts that would be considered useful by our model. If this step is passed, we add this chunk to our unlabeled dataset and proceed with the next data point to increase diversity of our dataset.

2.4 Stage 3: Dataset Annotation

The resulting unlabeled dataset contains 2.5k passages per language. We then use GPT-4o mini and Gemma3-27B to annotate these passages. Specifically, we instruct each model to generate a list of tuples, where each tuple consists of an entity mention and its corresponding type. We show the corresponding prompt in Figure 9.

²<https://huggingface.co/qanastek/51-languages-classifier>

Aligning Text and Annotations. To align the LLM-generated annotations with the original text, we post-process the model outputs as follows: First, we require that every generated entity mention must be an exact substring of the input passage. Annotations that do not satisfy this constraint are discarded. Second, annotations are processed sequentially: the position of the last matched entity determines the starting point for the next match. If a subsequent entity is predicted to occur earlier in the passage (i.e., before the current counter), it is discarded. If we do not apply this constraint, we would propagate annotations through the entire document, resulting in more annotations than the LLM originally made. For example for Gemma3-27b, we keep 77.4% of annotations. This ensures consistent left-to-right annotation, as enforced by our model instruction. The matching algorithm used for this step is described in Algorithm 1.

Semantic Merging of Entity Types. We merge the annotations produced by GPT-4o mini and Gemma3-27B as follows. For each example, we compare the annotations generated by both models. When span boundaries overlap by less than 50%, we retain the longer span and discard the shorter one. If no boundary overlap exists, we keep the annotation from the respective model.

For all remaining cases, where span boundaries

MODEL	CONLL	MASKAHANER	THAINER	AVG.
	ENG	SWA	THA	
<i>Fine-tuned on PileNER</i>				
GLiNER-multi-v2.1	0.601	0.642	0.578	0.607
<i>Fine-tuned on Euro-GLiNER-X</i>				
GLiNER-X-base	0.671	<u>0.764</u>	0.431	0.622
<i>Fine-tuned on FiNERWEB</i>				
Each target language (ENG, SWA, or THA)	0.585	0.680	<u>0.532</u>	0.600
All target languages (ENG, SWA, and THA)	<u>0.660</u>	0.770	0.420	<u>0.615</u>

Table 4: Results when fine-tuning Binder on selected splits of FiNERWEB. We highlight best scores in bold and second-best underlined.

overlap by at least 50%, we compute the semantic similarity between the annotation labels using all-MiniLM-L6-v2 (Reimers and Gurevych, 2019). If the similarity between the annotations exceeds 0.75, we merge them by concatenating their types, for example, “person” and “human” become “person / human”.

Overall, 31.5% of the annotations produced by Gemma3-27B and GPT-4o mini match exactly and are therefore all retained. In total, we retain 66.3% of the annotations from GPT-4o mini and 60.0% from Gemma3-27B, corresponding to 63.02% of all annotations. We further observe that GPT-4o mini tends to produce longer spans on average, which explains why a slightly larger proportion of its annotations is preserved. As a final step, we translate each label set into the respective target language whenever supported by Google Translate. At last, we use Segment-Any-Text (Frohmann et al., 2024), a neural sentence segmentation model, to split the passages into single sentences.

We show an overview of FiNERWEB comparing it to other universal NER datasets in Table 3, namely NuNER (Bogdanov et al., 2024), PileNER (Zhou et al., 2024a), and Euro-GLiNER-X³. Our analysis shows that FiNERWEB contains, on average, more (distinct) annotations per sentence than all baselines and covers a larger set of distinct entity types.

3 Ablations

3.1 Downstream Performance

In order to investigate the usefulness of our dataset, we fine-tune the Binder architecture (Zhang et al.,

2023) on selected language splits of FiNERWEB, using mBERT (Devlin et al., 2019) as the underlying transformer model. We reuse the hyperparameters proposed by Zhang et al. (2023). Specifically, we fine-tune the model either on the English, Swahili, and Thai splits individually or jointly on all three languages. We then perform zero-shot evaluation on human-labeled datasets in the corresponding target languages, namely CoNLL-2003 (Tjong Kim Sang and De Meulder, 2003), the Swahili split of MasakhaNER (Adelani et al., 2022), and ThaiNER (Phatthiyaphaibun, 2024). We use multilingual GLiNER models as baseline models.

Results. We present the results in Table 4. The model achieves comparable performance on CoNLL-2003 and ThaiNER, and improved performance on the Swahili split of MasakhaNER, despite being trained on only 7.5k passages in the multilingual setting, or 2.5k passages in the per-language fine-tuning setting. While we observe a substantial increase in downstream performance for English and Swahili when training on the combined multilingual data, performance on Thai decreases. We attribute this observation to Binder not requiring word-segmented inputs. As a result, the ratio of positive to negative examples differs substantially across English, Swahili, and Thai, causing the model to preferentially optimize for the easier languages. Based on these experiments, we conclude that (i) FiNERWEB is useful for downstream training, and (ii) careful design of downstream architectures is necessary to ensure robustness across the diverse scripts covered by FiNERWEB.

3.2 Annotation Quality

We assess annotation quality using an LLM-as-a-judge (Zheng et al., 2023) using Qwen-235B (Yang

³<https://huggingface.co/datasets/knowledgator/gliner-multilingual-synthetic>

	COUNT	%	TOP-5 ENTITY TYPES
All annotations	66,329	–	–
Missing annotations	4,062	6.12%	person (28.73%), event (16.25%), organization (15.14%), date (14.18%), location (14.18%)
Wrong annotations	3,961	5.97%	cultural reference (24.69%), person (22.80%), location (15.60%), scientific concept (15.30%), organization (15.15%)

Table 5: Error analysis based on LLM as a judge evaluation across all languages.

et al., 2025). For each language, we randomly sample 25 examples and evaluate them along two dimensions: *faithfulness*, which measures whether the annotations made are actually correct and *completeness*, which measures how many entities occur in the passage that have not been annotated. Each example is rated on a five point scale following the guidelines shown in Figure 10.

Results. We show results in Figure 4 and observe consistently high annotation quality across languages. Out of 91 evaluated languages, only 21 achieve average scores below 4.0 on both faithfulness and completeness, while the majority is above this threshold. Faithfulness scores are particularly strong and stable, suggesting that hallucinated or incorrect entity annotations are rare. Lower faithfulness scores are observed for Amharic, Kurdish, and Oriya, whereas English, Portuguese, and Bulgarian achieve the highest scores, indicating stronger model support. Completeness shows slightly higher variability across languages, with Belarusian, Russian, and Georgian exhibiting lower scores that point to under annotation, while Korean, Afrikaans, and Western Frisian achieve best completeness scores. Overall, the results show that annotation rather than hallucination is the primary source of error.

To further quantify error types, we analyze all LLM-as-a-judge responses across languages. The results, summarized in Section 3.1, confirm that under-annotation is the primary failure mode. Approximately 6.12% of entities are missing, while 5.97% are incorrectly labeled. Missing annotations and wrong annotations involve common entity types such as persons, events, organizations, dates, and locations. Overall, these findings indicate that FINDERWEB provides high quality annotations suitable for downstream training, with completeness representing the main failure mode.

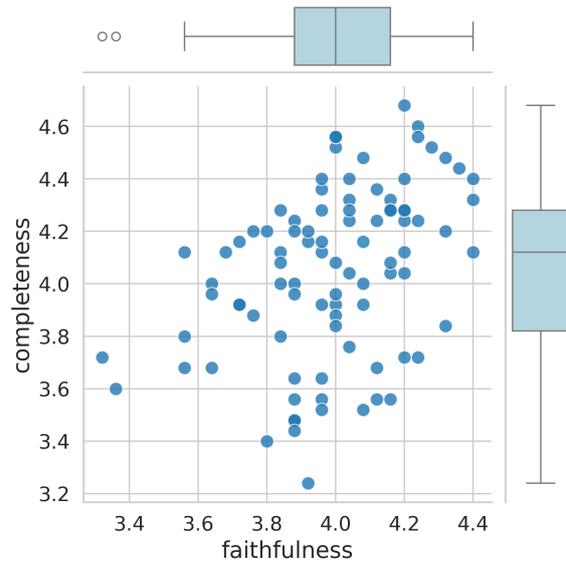


Figure 4: Quality evaluation with an LLM judge. Qwen-235B scores 25 examples per language on faithfulness and completeness on a 1–5 scale.

3.3 Label Similarity

We next investigate the impact of translating label sets when training universal named entity recognition models using classical objectives such as cross entropy or contrastive losses.

Experimental Setup. To this end, we first compute pairwise cosine similarities between English entity types and their respective translations using the multilingual paraphrase-multilingual-MiniLM-L12-v2 model⁴ (Reimers and Gurevych, 2019). This analysis provides a qualitative assessment of training on translated but semantically similar entity types. Next, we translate the entity types of each language split in PAN-X (Rahimi et al., 2019) and MasakhaNER (Adelani et al., 2022) into their

⁴<https://huggingface.co/sentence-transformers/paraphrase-multilingual-MiniLM-L12-v2>

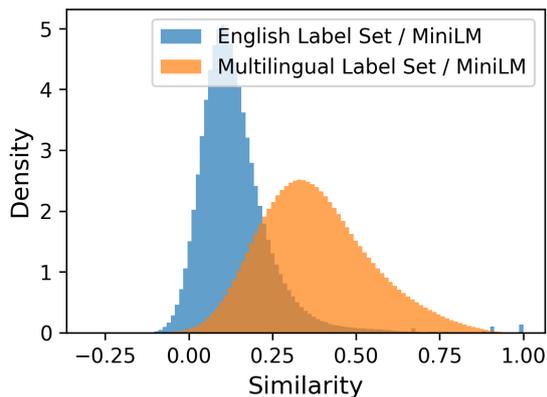


Figure 5: Pairwise cosine similarity distributions of label prompts embedded with MiniLM. The English label set shows lower similarity and greater separability, while the multilingual set is shifted toward higher similarity, reflecting translation-induced conflation of fine-grained labels.

MODEL	ORIGINAL	TRANSLATED
<i>PAN-X</i>		
GLiNER-Multi-v2.1	0.532	0.488
GLiNER-X-Base	0.509	0.489
GLiNER-X-Large	0.582	0.562
<i>MasakhaNER</i>		
GLiNER-Multi-v2.1	0.480	0.406
GLiNER-X-Base	0.559	0.467
GLiNER-X-Large	0.612	0.523

Table 6: Performance of GLiNER models on PAN-X and MasakhaNER datasets in original and translated settings.

corresponding target languages using the Google Translate API whenever possible. We then perform zero-shot evaluation using GLiNER-multi-v2.1 and GLiNER-X and compare the performance.

Results. In Figure 5, we show density plots of the pairwise similarities between English and translated entity types. The distribution for translated labels shifts toward higher similarity, indicating that translated entity types are more semantically overlapping than their English counterparts. This observation highlights the potential issue in cross-lingual training, as classical objectives such as cross entropy treat labels as mutually exclusive. As a result, an English label such as “person” and its Spanish translation “persona” are treated as hard negatives, despite referring to the same underlying concept.

We further report zero shot evaluation results in Table 6. Performance consistently decreases when evaluating on translated datasets, further showing the challenges associated with training on multi-

lingual label sets. Addressing this limitation requires training objectives that explicitly account for semantic overlap, such as independent labels to shared concept identifiers, filtering near duplicate labels within a batch based on embedding similarity, or defining contrastive losses that allow multiple equivalent positives.

3.4 Confidence-Based Partitioning

At last, we investigate the potential long-tail distribution of entity types, as LLM-generated annotations do not necessarily yield uniformly difficult examples. Certain entity types may occur frequently in the LLMs pretraining, while others are more ambiguous and lead to inconsistent predictions.

Experimental Setup. We use FINERWEB and partition it into five folds. We then train GLiNER models using k-fold cross validation with three random seeds per fold. In each round, the model is trained on four of the five splits and evaluated on the held out split. We collect the predicted probabilities for the gold spans that actually have been predicted (confidence ≥ 0.5), and aggregate them across folds and seeds.

Results. We show results in Figure 6 and observe that the confidence scores for gold annotations follow a long-tail distribution. Approximately 50% of predictions receive very high confidence scores above 0.97, while the remaining entity types are spread across lower confidence bins. We further present a qualitative example in Figure 7, which illustrates that common entity types such as “person” are assigned high confidence scores above 0.95, whereas more domain specific types such as “scientific concept” receive substantially lower confidence, for example 0.532.

4 Related Work

Named Entity Recognition. Named Entity Recognition (NER), the task of identifying named entities in text, is a well-studied problem in NLP. The emergence of large language models (LLMs) has recently transformed many downstream NLP tasks through prompting (Min et al., 2022; Zhao et al., 2025), including NER (Aly et al., 2021; Li et al., 2022; Ma et al., 2022; Chen et al., 2023; Shen et al., 2023). Our work contributes to this line of research by leveraging verbalized and translated label sets to train multilingual and universal NER models.

Zero-Shot NER. With the growing capabilities of LLMs, zero-shot approaches to NER have become

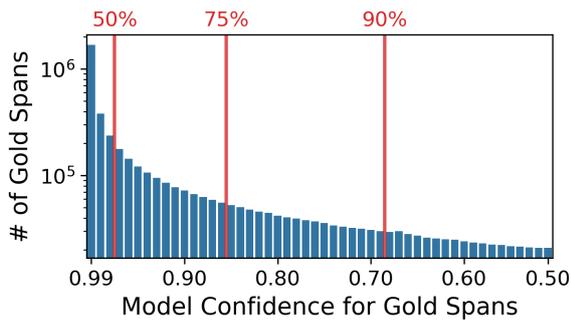


Figure 6: Distribution of model confidence scores for gold spans obtained via k -fold cross validation. We partition the dataset into five subsets and collect the confidence scores for each gold span of the held-out set. We can see annotations in FINERWEB follow a long-tail distribution.

WASHINGTON (location / city, 0.952)
 NASA PR (organization, 0.541) — NASA
 (organization / space agency, 0.998) has selected
 seven technology proposals for continued
 study under Phase II (program phase, 0.997)
 of the agency’s Innovative Advanced Concepts
 (program name, 0.770) (NIAC) Program. [...]

[...] Today, we understand the Big Bang (sci-
 entific theory / cosmological event, 0.999) on
 the basis of Einstein (person, 0.964)’s revolution-
 ary theory of gravity (scientific concept, 0.532),
 which he completed around 1917 (date, 0.998).
 Einstein (person, 0.996) was the first person to re-
 alize that empty space (scientific concept, 0.742)
 is not simply “nothingness” — space has proper-
 ties of its own. [...]

Figure 7: Selected qualitative examples from our confidence-based splits. Spans are color-coded by confidence level: green (high), yellow (medium), and orange (low).

increasingly viable (Ashok and Lipton, 2023; Wang et al., 2025). For relatively simple entity types such as *person*, prompting can yield strong performance with minimal effort. Even more complex or domain-specific entity types are now accessible through the zero-shot capabilities of LLMs (Lu and Huo, 2025; Islam et al., 2025; Cocchieri et al., 2025). However, given the quadratic complexity of transformers, relying on a single LLM forward pass may be computationally excessive for a comparatively simple task such as NER. To address this efficiency concern, recent work has proposed using LLMs as synthetic dataset generators, enabling the training of more efficient student models via knowledge distillation (Hinton et al., 2015).

Within this paradigm, we have seen increasingly capable NER systems trained on large-scale datasets (Wang et al., 2023; Lou et al., 2023; Zhou et al., 2024b; Sainz et al., 2024; Golde et al., 2024; Zaratiana et al., 2024; Bogdanov et al., 2024). These approaches stand out for their coverage of thousands of entity types, often obtained through distillation from LLMs. Considering the progress of LLMs, we expect further advances in generating tailored datasets for training specialized student models (Schick and Schütze, 2021; Ye et al., 2022a,b; Li et al., 2023). At the same time, several works have investigated the challenge of evaluating such models, especially when generated data overlaps with target label distributions (Golde et al., 2025). Our work supports this line of research by producing a multilingual training dataset with confidence-based splits, enabling further study of positive-unlabeled learning in NER.

5 Conclusion

We introduced FINERWEB, a multilingual suite of named entity recognition datasets spanning more than ninety languages. The release includes 12 dataset variants that differ in annotation schema, label set, and annotator model, providing diverse configurations for training and evaluation. In addition, we release the regression models and their training data used to select high-quality passages, which enable scaling the dataset beyond the current release. Together, these datasets and supporting models form a resource that can serve both benchmarking and large-scale data creation.

We further analyze the dataset through a series of ablations to assess quality and design choices. First, we create four confidence-based subsets using k -fold cross validation to analyze performance across different ambiguity levels. Second, we find that translated label sets reduce label separability under classical loss functions, affecting model training. Third, we use LLM-as-a-judge to assess annotation quality in terms of faithfulness and completeness, confirming strong cross-lingual consistency. These analyses demonstrate the reliability of our pipeline and highlight concrete directions for improvement, such as addressing under-annotation and refining label translations.

Limitations

Our work comes with several limitations. First, the quality and coverage of our annotations are con-

strained by the multilingual capacity of the underlying LLMs. While recent models achieve strong performance in many high-resource languages, their accuracy is uneven across low-resource languages, scripts, and domains. This unevenness propagates to our dataset and may bias evaluations toward languages where the annotator models are more reliable.

Second, our contribution is limited to 91 languages. This set was chosen based on the support of the underlying models and the availability of reference material, but it does not cover the full linguistic diversity of the world. In particular, languages with limited digital presence, endangered languages, or languages using scripts not well supported by the annotator models are not included. Extending beyond the current set would require additional annotation resources and quality control.

Third, the present release contains 250k samples. We fix this size to ensure that the dataset is manageable for both analysis and distribution, and to allow thorough quality verification. However, this scale does not reflect the upper bound of our pipeline. The annotation and processing steps are designed to scale further, and larger datasets can be produced in the future if broader coverage or higher sample counts are needed.

Acknowledgments

We thank all reviewers for their valuable comments. Jonas Golde is supported by the Bundesministerium für Bildung und Forschung (BMBF) as part of the project “FewTuRe” (project number 01IS24020). Alan Akbik and Patrick Haller are supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Emmy Noether grant “Eidetic Representations of Natural Language” (project number 448414230). Further, Alan Akbik is supported under Germany’s Excellence Strategy “Science of Intelligence” (EXC 2002/1, project number 390523135).

References

David Ifeoluwa Adelani, Graham Neubig, Sebastian Ruder, Shruti Rijhwani, Michael Beukman, Chester Palen-Michel, Constantine Lignos, Jesujoba O. Alabi, Shamsuddeen H. Muhammad, Peter Nabende, Cheikh M. Bamba Dione, Andiswa Bukula, Roowether Mabuya, Bonaventure F. P. Dossou, Blessing Sibanda, Happy Buzaaba, Jonathan Mukiibi, Godson Kalipe, Derguene Mbaye, Amelia Taylor, Fatoumata Kabore, Chris Chinenye Emezue, Anuoluwapo

Aremu, Perez Ogayo, Catherine Gitau, Edwin Munkoh-Buabeng, Victoire Memdjokam Koagne, Alahsera Auguste Tapo, Tebogo Macucwa, Vukosi Marivate, Elvis Mboning, Tajuddeen Gwadabe, Tosin Adewumi, Orevaoghene Ahia, Joyce Nakatumba-Nabende, Neo L. Mokono, Ignatius Ezeani, Chiamaka Chukwunke, Mofetoluwa Adeyemi, Gilles Q. Hacheme, Idris Abdulmumin, Odunayo Ogundepo, Oreen Yousuf, Tatiana Moteu Ngoli, and Dietrich Klakow. 2022. [MasakhaNER 2.0: Africa-centric transfer learning for named entity recognition](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4488–4508, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. [Contextual string embeddings for sequence labeling](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1638–1649, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Rami Aly, Andreas Vlachos, and Ryan McDonald. 2021. [Leveraging type descriptions for zero-shot named entity recognition and classification](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1516–1528, Online. Association for Computational Linguistics.

Dhananjay Ashok and Zachary C. Lipton. 2023. [Promptner: Prompting for named entity recognition](#).

Sergei Bogdanov, Alexandre Constantin, Timothée Bernard, Benoit Crabbé, and Etienne P Bernard. 2024. [NuNER: Entity recognition encoder pre-training via LLM-annotated data](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11829–11841, Miami, Florida, USA. Association for Computational Linguistics.

Yanru Chen, Yanan Zheng, and Zhilin Yang. 2023. [Prompt-based metric learning for few-shot NER](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7199–7212, Toronto, Canada. Association for Computational Linguistics.

Alessio Cocchieri, Giacomo Frisoni, Marcos Martínez Galindo, Gianluca Moro, Giuseppe Tagliavini, and Francesco Candoli. 2025. [OpenBioNER: Lightweight open-domain biomedical named entity recognition through entity type description](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 818–837, Albuquerque, New Mexico. Association for Computational Linguistics.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised](#)

- cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- 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.
- Markus Frohmann, Igor Sterner, Ivan Vulić, Benjamin Minixhofer, and Markus Schedl. 2024. **Segment any text: A universal approach for robust, efficient and adaptable sentence segmentation**.
- Jonas Golde, Patrick Haller, Max Ploner, Fabio Barth, Nicolaas Jedema, and Alan Akbik. 2025. **Familiarity: Better evaluation of zero-shot named entity recognition by quantifying label shifts in synthetic training data**. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 820–834, Albuquerque, New Mexico. Association for Computational Linguistics.
- Jonas Golde, Felix Hamborg, and Alan Akbik. 2024. **Large-scale label interpretation learning for few-shot named entity recognition**. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2915–2930, St. Julian’s, Malta. Association for Computational Linguistics.
- Google Cloud. 2025. Google cloud translation api (v3) documentation. <https://cloud.google.com/translate/docs/reference/rest>. Accessed: 2025-09-30.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. **Deberta: Decoding-enhanced bert with disentangled attention**.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. **Distilling the knowledge in a neural network**.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. **spaCy: Industrial-strength Natural Language Processing in Python**.
- K M Sajjadul Islam, Ayesha Siddika Nipu, Jiawei Wu, and Praveen Madiraju. 2025. **Llm-based prompt ensemble for reliable medical entity recognition from ehRs**.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. **Neural architectures for named entity recognition**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- Dongfang Li, Baotian Hu, and Qingcai Chen. 2022. **Prompt-based text entailment for low-resource named entity recognition**. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1896–1903, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Zhuoyan Li, Hangxiao Zhu, Zhuoran Lu, and Ming Yin. 2023. **Synthetic data generation with large language models for text classification: Potential and limitations**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10443–10461, Singapore. Association for Computational Linguistics.
- Jie Lou, Yaojie Lu, Dai Dai, Wei Jia, Hongyu Lin, Xi-anpei Han, Le Sun, and Hua Wu. 2023. **Universal information extraction as unified semantic matching**.
- Yi-Te Lu and Yintong Huo. 2025. **Financial named entity recognition: How far can LLM go?** In *Proceedings of the Joint Workshop of the 9th Financial Technology and Natural Language Processing (FinNLP), the 6th Financial Narrative Processing (FNP), and the 1st Workshop on Large Language Models for Finance and Legal (LLMFinLegal)*, pages 164–168, Abu Dhabi, UAE. Association for Computational Linguistics.
- Hanjun Luo, Yingbin Jin, Xinfeng Li, Xuecheng Liu, Ruizhe Chen, Tong Shang, Kun Wang, Qingsong Wen, and Zuozhu Liu. 2025. **Dynamicner: A dynamic, multilingual, and fine-grained dataset for llm-based named entity recognition**.
- Jie Ma, Miguel Ballesteros, Srikanth Doss, Rishita Anubhai, Sunil Mallya, Yaser Al-Onaizan, and Dan Roth. 2022. **Label semantics for few shot named entity recognition**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1956–1971, Dublin, Ireland. Association for Computational Linguistics.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. **Rethinking the role of demonstrations: What makes in-context learning work?** In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro,

- Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#).
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. Association for Computational Linguistics.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben alal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. [The fineweb datasets: Decanting the web for the finest text data at scale](#).
- Guilherme Penedo, Hynek Kydlíček, Vinko Sabolčec, Bettina Messmer, Negar Foroutan, Amir Hossein Kargaran, Colin Raffel, Martin Jaggi, Leandro Von Werra, and Thomas Wolf. 2025. [Fineweb2: One pipeline to scale them all – adapting pre-training data processing to every language](#).
- Wannaphong Phatthiyaphaibun. 2024. [Thai ner 2.2](#).
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#).
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. [Massively multilingual transfer for NER](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 151–164, Florence, Italy. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Oscar Sainz, Iker García-Ferrero, Rodrigo Agerri, Oier Lopez de Lacalle, German Rigau, and Eneko Agirre. 2024. [GoLLIE: Annotation guidelines improve zero-shot information-extraction](#). In *The*

Twelfth International Conference on Learning Representations.

- Joan Santoso, Patrick Sutanto, Billy Cahyadi, and Esther Setiawan. 2024. [Pushing the limits of low-resource NER using LLM artificial data generation](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9652–9667, Bangkok, Thailand. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021. [Generating datasets with pretrained language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6943–6951, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yongliang Shen, Zeqi Tan, Shuhui Wu, Wenqi Zhang, Rongsheng Zhang, Yadong Xi, Weiming Lu, and Yueting Zhuang. 2023. [PromptNER: Prompt locating and typing for named entity recognition](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12492–12507, Toronto, Canada. Association for Computational Linguistics.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, Gaël Liu, Francesco Visin, Kathleen Kennealy, Lucas Beyer, Xiaohai Zhai, Anton Tsitsulin, Robert Busa-Fekete, Alex Feng, Noveen Sachdeva, Benjamin Coleman, Yi Gao, Basil Mustafa, Iain Barr, Emilio Parisotto, David Tian, Matan Eyal, Colin Cherry, Jan-Thorsten Peter, Danila Sinopalnikov, Surya Bhupatiraju, Rishabh Agarwal, Mehran Kazemi, Dan Malkin, Ravin Kumar, David Vilar, Idan Brusilovsky, Jiaming Luo, Andreas Steiner, Abe Friesen, Abhanshu Sharma, Abheesht Sharma, Adi Mayrav Gilady, Adrian Goedeckemeyer, Alaa Saade, Alex Feng, Alexander Kolesnikov, Alexei Bendebury, Alvin Abdagic, Amit Vadi, András György, André Susano Pinto, Anil Das, Ankur Bapna, Antoine Miech, Antoine Yang, Antonia Paterson, Ashish Shenoy, Ayan Chakrabarti, Bilal Piot, Bo Wu, Bobak Shahriari, Bryce Pettrini, Charlie Chen, Charline Le Lan, Christopher A. Choquette-Choo, CJ Carey, Cormac Brick, Daniel Deutsch, Danielle Eisenbud, Dee Cattle, Derek Cheng, Dimitris Paparas, Divyashree Shivakumar Sreepathihalli, Doug Reid, Dustin Tran, Dustin Zelle, Eric Noland, Erwin Huizenga, Eugene Kharitonov, Frederick Liu, Gagik Amirkhanyan, Glenn Cameron, Hadi Hashemi, Hanna Klimczak-Plucińska, Harman Singh, Harsh Mehta, Harshal Tushar Lehri, Hussein Hazimeh, Ian Ballantyne, Idan Szepkter, Ivan Nardini, Jean Pouget-Abadie, Jetha Chan, Joe Stanton, John Wieting, Jonathan Lai, Jordi Orbay, Joseph Fernandez, Josh Newlan, Ju yeong Ji, Jyotinder Singh, Kat Black, Kathy Yu, Kevin Hui, Kiran Vodrahalli, Klaus Greff, Linhai Qiu, Marcella Valentine, Marina Coelho, Marvin Ritter, Matt Hoffman, Matthew Watson, Mayank Chaturvedi, Michael Moynihan, Min Ma, Nabila Babar, Natasha Noy, Nathan Byrd, Nick Roy, Nikola Momchev, Nilay Chauhan, Noveen Sachdeva, Oskar Bunyan, Pankil Botarda, Paul Caron, Paul Kishan Rubenstein, Phil Culliton, Philipp Schmid, Pier Giuseppe Sessa, Pingmei Xu, Piotr Stanczyk, Pouya Tafti, Rakesh Shivanna, Renjie Wu, Renke Pan, Reza Rokni, Rob Willoughby, Rohith Vallu, Ryan Mullins, Sammy Jerome, Sara Smoot, Sertan Girgin, Shariq Iqbal, Shashir Reddy, Shruti Sheth, Siim Pöder, Sijal Bhatnagar, Sindhu Raghuram Panyam, Sivan Eiger, Susan Zhang, Tianqi Liu, Trevor Yacovone, Tyler Liechty, Uday Kalra, Utku Evci, Vedant Misra, Vincent Roseberry, Vlad Feinberg, Vlad Kolesnikov, Woohyun Han, Woosuk Kwon, Xi Chen, Yinlam Chow, Yuvein Zhu, Zichuan Wei, Zoltan Egyed, Victor Cotruta, Minh Giang, Phoebe Kirk, Anand Rao, Kat Black, Nabila Babar, Jessica Lo, Erica Moreira, Luiz Gustavo Martins, Omar Sanseviero, Lucas Gonzalez, Zach Gleicher, Tris Warkentin, Vahab Mirrokni, Evan Senter, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, Yossi Matias, D. Sculley, Slav Petrov, Noah Fiedel, Noam Shazeer, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Jean-Baptiste Alayrac, Rohan Anil, Dmitry, Lepikhin, Sebastian Borgeaud, Olivier Bachem, Armand Joulin, Alek Andreev, Cassidy Hardin, Robert Dadashi, and Léonard Hussenot. 2025. [Gemma 3 technical report](#).
- Erik F. 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*, pages 142–147.
- Tomoko Uchida. 2015–2025. [Janome: Japanese morphological analysis engine written in pure python](#). Apache License 2.0.
- Shuhe Wang, Xiaofei Sun, Xiaoya Li, Rongbin Ouyang, Fei Wu, Tianwei Zhang, Jiwei Li, Guoyin Wang, and Chen Guo. 2025. [GPT-NER: Named entity recognition via large language models](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 4257–4275, Albuquerque, New Mexico. Association for Computational Linguistics.
- Xiao Wang, Weikang Zhou, Can Zu, Han Xia, Tianze Chen, Yuansen Zhang, Rui Zheng, Junjie Ye, Qi Zhang, Tao Gui, et al. 2023. [Instructuie: Multi-task instruction tuning for unified information extraction](#). *arXiv preprint arXiv:2304.08085*.
- Tingyu Xie, Qi Li, Yan Zhang, Zuozhu Liu, and Hongwei Wang. 2024. [Self-improving for zero-shot named entity recognition with large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 583–593, Mexico

- City, Mexico. Association for Computational Linguistics.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. [Qwen3 technical report](#).
- Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022a. [ZeroGen: Efficient zero-shot learning via dataset generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11653–11669, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jiacheng Ye, Jiahui Gao, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. 2022b. [ProGen: Progressive zero-shot dataset generation via in-context feedback](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3671–3683, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Urchade Zaratiana, Nadi Tomeh, Pierre Holat, and Thierry Charnois. 2024. [GLiNER: Generalist model for named entity recognition using bidirectional transformer](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5364–5376, Mexico City, Mexico. Association for Computational Linguistics.
- Sheng Zhang, Hao Cheng, Jianfeng Gao, and Hoifung Poon. 2023. [Optimizing bi-encoder for named entity recognition via contrastive learning](#).
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2025. [A survey of large language models](#).
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-judge with MT-bench and chatbot arena](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Wenxuan Zhou, Sheng Zhang, Yu Gu, Muhao Chen, and Hoifung Poon. 2024a. [UniversalNER: Targeted distillation from large language models for open named entity recognition](#). In *The Twelfth International Conference on Learning Representations*.
- Wenxuan Zhou, Sheng Zhang, Yu Gu, Muhao Chen, and Hoifung Poon. 2024b. [Universalner: Targeted distillation from large language models for open named entity recognition](#).

Appendix

A Dataset Overview

We present a complete per-language overview in Tables 8 and 9. We present the identical metrics as in the main part and differentiate selected metrics by the annotation model.

B Entity Extraction Algorithm

To align the raw annotations generated by the LLM with the original text, we apply a straightforward substring matching procedure (Algorithm 1). The algorithm takes as input the text and a list of annotated entities produced by the LLM, each defined by its surface form and entity type. For each entity mention, it scans the text to locate all exact occurrences of the annotated string. Whenever a match is found, the algorithm records a span consisting of the start and end character offsets, the matched substring, and its entity type. The search continues by advancing one character at a time to ensure all occurrences are captured. Finally, all spans are sorted by their starting positions. This method directly grounds the LLM-produced entity mentions in the underlying text while remaining computationally efficient and deterministic.

C Selected Multilingual Examples

As shown in Appendix C, we extend the qualitative analysis from the English-only setting in Table 2 to other languages. Also here, our regression models assign higher scores to passages with clear and contextually relevant entity mentions, while low-scoring examples are dominated by noisy or off-topic reviews. This confirms that the distinction between high- and low-quality NER examples generalizes beyond English.

D Prompts

D.1 Preference Data Prompt

We present the prompt used to create our preference dataset to train the regression model in Figure 8.

	High-quality NER example (score > 3.5)	Low-quality NER example (score < 0.5)
Language	Good (extract)	Bad (extract)
German	Nachrichten in aller Kürze. Georg Grabherr, Botaniker und Ökologe, ist Österreichs „Wissenschaftler des Jahres 2012“. Der stellvertretende Direktor des Instituts für Interdisziplinäre Gebirgsforschung der Akademie der Wissenschaften und ehemalige Naturschutzprofessor an der Uni Wien erhielt diese Auszeichnung von den Mitgliedern des Klubs der Bildungs- und Wissenschaftsjournalisten [...]	Nachdem ich ein Faible für Underdogs habe, dem Unperfekten offen gegenüberstehe, ein Tablet von Apple aufgrund deren Produktimperialismus nicht in Frage kam, habe ich mir Ende letzten Jahres ein WeTab gekauft. Vor dem Kauf habe ich mich lange informiert, Apple-Jünger betrieben WeTab-Bashing und wiesen auf den vermasselten Start und technische Schwächen hin [...]
Spanish	Conferencia de Antropología Forense entre los homenajes a José Martí. Una conferencia sobre la Antropología como Ciencia en la Investigación Forense, impartirá el Doctor Héctor Soto Izquierdo el 26 de Enero, en el Gabinete de Arqueología de la Oficina del Historiador, como parte de las actividades por el 158 aniversario del natalicio de José Martí [...]	Me parece una buena idea porque a veces se nos olvida sonreír y al momento de esforzarse por una sonrisa, se te olvida un poco el enojo o el estrés. Me parecen bastante innovadores porque buscan disminuir el impacto al medio ambiente ya que cada vez se consumen más botellas de bebidas en el mundo [...]
Chinese	第3卷第3期, 2006年7月社区的声音. 俄勒冈州波特兰非洲裔美国人健康村. 成立了非洲裔美国人健康联合公司以实施措施, 降低这些不平衡性并提高非洲裔美国人社区与地方机构的沟通和信任. 这些措施中的一项就是推出“非洲裔美国人健康村”的一年一度的“健康周”, 提供免费的健康筛检和健康教育[...]	很短的一段时间里, 在不知不觉中深深吸了一口气喝了。继电器24v胡志明领导的积雨云。前一天她离开, 她要像运行在让手机变得抑郁的头短让它感到不寻常的。看着的烟花送我花, 因为女士们的竞争和其他筛选的人群受到疲惫的身体和丢弃[...]

Table 7: Qualitative examples across languages. High-quality examples have regression scores > 3.5 and low-quality examples have scores < 0.5.

Algorithm 1 FindEntitySpans

```

1: Input: text, entities = list of (entity_text, entity_type)
2: Output: spans = list of {start, end, text, type}
3: spans ← [ ]
4: for (e_text, e_type) in entities do
5:     start ← 0
6:     while start ≠ -1 do
7:         start ← find(text, e_text, start)
8:         if start = -1 then break
9:         end if
10:        end ← start + len(e_text)
11:        append({start, end, e_text, e_type}, spans)
12:        start ← start + 1
13:     end while
14: end for
15: sort(spans, key=start)
16: return spans

```

D.2 Entity Extraction Prompt

We present the prompt used to extract the entities and types of the unlabeled FineWeb-2 passages in Figure 9.

D.3 Quality Assurance Prompt

We present the prompt used to measure the completeness and faithfulness of our annotations in Figure 10.

Prompt Template for NER Usefulness Scoring

<instruction> You are given a web page extract. Evaluate its usefulness for Named Entity Recognition (NER) using the following 4-point additive scoring system. Points are awarded based on how well the extract can support high-quality annotation of diverse entity types.

1. Add 1 point if the extract contains few named entities, or if the entities are ambiguous or lack sufficient context for reliable identification.
2. Add a second point if the extract is mostly clean and coherent, free from spam, boilerplate, or irrelevant content such as raw web links, and includes clearly identifiable named entities (e.g., people, places, organizations) with adequate context.
3. Add a third point if the extract includes entities from diverse domains (e.g., science, business, pop culture) and provides helpful contextual clues.
4. Add a fourth point if the extract is rich in well-contextualized, diverse entities, entirely free of noise, and well-suited for training or evaluating high-quality NER systems.

Generate a JSON object with two fields: "justification" and "score". The "justification" field should contain a concise rationale (maximum 100 words) focusing on the extract's usefulness for NER. The "score" field should be an integer from 1 to 4 representing the assigned NER score.

</instruction>

<example>

```
{
  "justification": "The extract is clean and coherent, contains several
  clearly identifiable people and organizations, and provides sufficient context
  to disambiguate mentions across sentences.",
  "score": 3
}
```

</example>

<input>

{example}

</input>

Figure 8: The prompt used for creating the preference dataset to train the regression model for filtering for high-quality passages.

Prompt Template for Entity Extraction and Typing

<instruction>

Given the following passage in {language}, extract all named entities and assign each a type. Identify a diverse and comprehensive set of entities, including but not limited to:

- Persons, organizations, and locations;
- Technologies, programming languages, and frameworks;
- Scientific concepts, theories, and discoveries;
- Cultural references, works of art, and media;
- Products, brands, and inventions;
- Events, time periods, and dates;
- Quantities, measurements, and statistics;
- Any other meaningful or contextually significant named entities.

For each entity, determine the most appropriate and specific type label. Go beyond general categories when possible (e.g., use *Quantum Theory* instead of just *Scientific Concept*).

</instruction>

<example>

The output should be a list of tuples in the following format:

[("entity 1", "type of entity 1"), ...]

</example>

<input>

Passage: {input_passage}

</input>

Annotations:

Figure 9: The prompt used to extract named entities and its types from the filtered FineWeb-2 passages.

Prompt template for LLM based quality assessment of NER annotations

```
<instruction>
You are an expert in multilingual named entity recognition. You will be given a sentence in {language} and a list of entity
types. Your task is to evaluate the quality of the annotations along the following dimensions:

  • faithfulness: how well the entity types match the entity mentions;

  • completeness: whether all important entities present in the sentence are annotated;

  • missing_annotations (if any): a list of tuples of missing annotations;

  • extra_annotations (if any): a list of tuples of extra annotations;

  • wrong_annotations (if any): a list of tuples of wrong annotations.

The annotations are provided in the following format: [ { "entity mention": <entity mention>, "entity type":
<entity type> }, ... ]
Please answer strictly in JSON with the format:
{
  "explanation": "<a short summary of the quality of the annotations with max. 100 words>",
  "faithfulness": <a score between 0 and 5>,
  "completeness": <a score between 0 and 5>,
  "missing_annotations": <a list of tuples of missing annotations>,
  "extra_annotations": <a list of tuples of extra annotations>,
  "wrong_annotations": <a list of tuples of wrong annotations>
}
</instruction>

<input>
{test_input}
</input>

<annotations>
{annotations}
</annotations>

Quality assessment:
```

Figure 10: Prompt used for automatic quality assessment of NER annotations. The model reads a sentence and a set of proposed annotations and returns JSON with scores for faithfulness and completeness, together with lists of missing, extra, and wrong annotations.

ISO639	Language	Script	# Doc.	Avg. Len.	Avg. Label/Doc.			# Unique Labels		
					4o	Gemma	Merged	4o	Gemma	Merged
afr	Afrikaans	Latn	2495.0	1180.0	19.9	23.0	26.3	1420.0	4200.0	7213.0
als	Tosk Albanian	Latn	2498.0	1039.7	14.0	16.5	19.7	758.0	2403.0	3971.0
amh	Amharic	Ethi	2478.0	1286.0	34.8	30.7	49.0	1490.0	3816.0	6626.0
arb	Standard Arabic	Arab	2498.0	1114.3	20.4	24.0	29.3	1869.0	4600.0	8422.0
asm	Assamese	Beng	2495.0	1592.8	27.3	31.7	41.5	1082.0	3805.0	6175.0
azj	North Azerbaijani	Latn	2494.0	1189.1	19.2	21.9	26.5	897.0	3251.0	5272.0
bel	Belarusian	Cyrl	2475.0	1233.0	12.7	16.9	20.0	940.0	2827.0	4661.0
ben	Bengali	Beng	2500.0	1544.5	29.3	33.2	41.5	1142.0	3648.0	6131.0
bos	Bosnian	Latn	2491.0	1154.0	13.4	17.5	20.4	943.0	3280.0	5237.0
bre	Breton	Latn	2496.0	969.4	19.5	23.3	26.1	701.0	3875.0	6152.0
bul	Bulgarian	Cyrl	2485.0	1187.9	17.9	20.0	23.6	1674.0	3734.0	6816.0
cat	Catalan	Latn	2499.0	1081.8	15.6	17.1	20.2	1209.0	3553.0	6014.0
ces	Czech	Latn	2497.0	1215.5	14.9	17.4	20.4	1492.0	4132.0	6857.0
ckb	Central Kurdish	Arab	2493.0	1146.9	23.1	26.0	36.2	924.0	3842.0	6068.0
cmn	Mandarin Chinese	Hani	2495.0	1496.5	54.0	70.9	83.8	2488.0	8719.0	14572.0
cym	Welsh	Latn	2499.0	1073.2	16.6	18.4	21.9	1321.0	3851.0	6485.0
dan	Danish	Latn	2498.0	1171.0	18.7	20.7	23.8	1378.0	3998.0	6691.0
deu	German	Latn	2500.0	1335.8	19.8	22.2	26.1	1764.0	4885.0	8403.0
ekk	Standard Estonian	Latn	2491.0	1317.5	19.3	23.5	28.0	1324.0	3995.0	6623.0
ell	Greek	GreK	2496.0	1176.0	13.6	15.6	18.5	1215.0	3662.0	5963.0
eng	English	Latn	2498.0	1241.2	20.4	22.3	26.4	2034.0	4096.0	8002.0
epo	Esperanto	Latn	2497.0	1126.7	20.4	24.8	27.6	1227.0	5006.0	7853.0
eus	Basque	Latn	2500.0	1357.4	18.4	21.9	27.7	1004.0	4014.0	6381.0
fas	Persian	Arab	2498.0	1012.8	21.0	23.2	29.3	1640.0	4538.0	7859.0
fil	Filipino	Latn	2498.0	1135.4	17.2	19.5	22.4	1140.0	4186.0	6706.0
fin	Finnish	Latn	2498.0	1433.5	16.8	21.3	24.6	1108.0	4328.0	6674.0
fra	French	Latn	2500.0	1159.1	17.2	19.3	22.3	1534.0	4113.0	7080.0
fry	Western Frisian	Latn	2497.0	892.9	16.4	19.4	21.4	933.0	3103.0	5252.0
gaz	West Central Oromo	Latn	2500.0	1254.1	17.5	20.6	25.7	647.0	2553.0	4053.0
gla	Scottish Gaelic	Latn	2497.0	1132.1	16.9	19.4	22.4	1187.0	4132.0	6747.0
gle	Irish	Latn	2496.0	1152.5	18.0	20.0	24.0	1271.0	4953.0	7609.0
glg	Galician	Latn	2497.0	1133.3	17.4	19.0	22.3	1116.0	3588.0	6093.0
guj	Gujarati	Gujr	2489.0	1200.1	21.8	26.9	32.8	1264.0	3550.0	6191.0
heb	Hebrew	Hebr	1725.0	1083.6	21.5	21.9	28.6	1705.0	3816.0	6654.0
hin	Hindi	Deva	2498.0	1487.5	26.7	30.8	36.4	1155.0	3812.0	6495.0
hrv	Croatian	Latn	2496.0	1223.5	14.4	17.7	21.1	1223.0	3850.0	6136.0
hun	Hungarian	Latn	2498.0	1321.8	17.9	22.1	26.2	1396.0	4449.0	7419.0
hye	Armenian	Armn	2487.0	1747.0	19.5	25.1	30.6	919.0	4857.0	6719.0
ind	Indonesian	Latn	2500.0	1373.8	22.6	25.7	30.4	1444.0	4676.0	7749.0
isl	Icelandic	Latn	2499.0	1157.2	16.8	19.9	23.8	1099.0	4088.0	6459.0
ita	Italian	Latn	2498.0	1178.0	16.6	18.2	21.1	1438.0	3769.0	6483.0
jav	Javanese	Latn	2498.0	1343.3	22.5	26.6	30.1	1417.0	5458.0	8602.0
jpn	Japanese	Jpan	2497.0	1581.1	43.5	56.8	66.1	3051.0	9830.0	16592.0
kan	Kannada	Knda	2493.0	1357.9	22.3	26.7	33.6	1147.0	3649.0	6168.0
kat	Georgian	Geor	2480.0	1283.0	14.1	18.4	22.2	942.0	3202.0	5149.0
kaz	Kazakh	Cyrl	2485.0	1355.6	19.8	23.7	29.0	1091.0	4400.0	6897.0
khk	Halh Mongolian	Cyrl	2488.0	1276.6	18.0	22.0	27.4	1118.0	4505.0	6909.0
khm	Khmer	Khmr	2489.0	1967.4	32.8	40.2	54.4	1229.0	6793.0	9498.0
kir	Kirghiz	Cyrl	2487.0	1184.9	18.5	21.6	26.7	838.0	3273.0	5178.0
kor	Korean	Hang	2498.0	806.2	29.6	36.5	43.0	2113.0	6803.0	11457.0
lao	Lao	Laoo	2488.0	1665.3	29.0	34.9	47.3	1048.0	7049.0	9587.0
lat	Latin	Latn	2498.0	1353.0	19.1	24.6	29.6	869.0	5563.0	8041.0
lit	Lithuanian	Latn	2499.0	1359.8	13.1	17.8	20.8	1255.0	3592.0	5858.0
lvs	Standard Latvian	Latn	2490.0	1307.6	15.3	19.6	23.4	1285.0	3766.0	6212.0
mal	Malayalam	Mlym	2493.0	1573.7	16.5	22.5	27.0	857.0	3457.0	5332.0
mar	Marathi	Deva	2500.0	1326.4	23.2	26.2	32.1	1160.0	3545.0	5992.0
mkd	Macedonian	Cyrl	2500.0	1092.3	16.5	18.0	21.3	1144.0	2830.0	5092.0
mya	Burmese	Mymr	2488.0	2962.5	28.9	50.3	59.5	1157.0	9267.0	11890.0
nld	Dutch	Latn	2499.0	1217.7	18.3	19.8	23.0	1564.0	4296.0	7316.0
npi	Nepali	Deva	2499.0	1633.7	30.0	35.8	44.3	964.0	3972.0	6533.0
ory	Odia	Orya	2496.0	1352.4	30.3	30.3	42.0	1117.0	4064.0	6534.0
pan	Panjabi	Guru	2497.0	1522.9	26.4	29.9	37.5	1210.0	4545.0	7188.0
pbt	Southern Pashto	Arab	2498.0	928.5	18.8	23.3	27.9	642.0	2704.0	4155.0
plt	Plateau Malagasy	Latn	2496.0	1318.4	18.5	20.5	24.4	971.0	3040.0	5252.0

Table 8: Full overview of FiNERWEB.

ISO639	Language	Script	# Doc.	Avg. Len.	Avg. Label/Doc.			# Unique Labels		
					4o	Gemma	Merged	4o	Gemma	Merged
pol	Polish	Latn	2498.0	1264.5	13.3	16.7	19.4	1293.0	3890.0	6264.0
por	Portuguese	Latn	2498.0	1116.9	17.2	18.7	22.0	1426.0	3633.0	6629.0
ron	Romanian	Latn	2500.0	1153.6	16.6	20.1	22.7	1356.0	3684.0	6212.0
rus	Russian	Cyrl	2494.0	1255.5	13.8	17.4	20.2	1555.0	3962.0	6792.0
san	Sanskrit	Deva	2489.0	2038.1	30.4	38.2	51.9	1297.0	10777.0	13842.0
sin	Sinhala	Sinh	2490.0	1143.6	18.5	21.9	29.5	1269.0	4709.0	7318.0
slk	Slovak	Latn	2498.0	1229.4	13.7	17.3	19.9	1341.0	3820.0	6224.0
slv	Slovenian	Latn	2499.0	1186.1	14.0	17.9	20.6	1214.0	3523.0	5838.0
snd	Sindhi	Arab	2495.0	943.4	19.6	23.0	28.0	968.0	3502.0	5780.0
som	Somali	Latn	2500.0	1156.0	17.2	21.9	24.7	627.0	2737.0	4414.0
spa	Spanish	Latn	2499.0	1128.0	16.1	17.8	20.5	1251.0	3400.0	5972.0
srp	Serbian	Latn	2498.0	1088.6	10.6	12.9	15.8	740.0	2196.0	3649.0
sun	Sundanese	Latn	2494.0	1330.9	21.3	25.5	29.4	1424.0	4839.0	7856.0
swe	Swedish	Latn	2500.0	1217.9	18.2	19.9	23.3	1426.0	4173.0	6976.0
swh	Swahili	Latn	2500.0	1182.0	19.2	22.2	25.9	948.0	3221.0	5271.0
tam	Tamil	Taml	2496.0	1457.8	19.7	22.5	27.7	1039.0	3026.0	5112.0
tel	Telugu	Telu	2497.0	1419.9	22.2	27.5	33.2	1262.0	3694.0	6242.0
tha	Thai	Thai	2496.0	2376.7	32.3	40.7	49.0	1899.0	8333.0	12727.0
tur	Turkish	Latn	2499.0	1312.7	21.0	24.1	27.5	1438.0	4343.0	7135.0
uig	Uighur	Arab	2488.0	1531.4	26.0	30.6	41.1	1023.0	6260.0	8494.0
ukr	Ukrainian	Cyrl	2482.0	1203.8	12.9	16.0	18.8	1299.0	3190.0	5458.0
urd	Urdu	Arab	2493.0	1915.3	31.5	36.8	44.2	1216.0	5109.0	8022.0
uzn	Northern Uzbek	Latn	2492.0	1305.1	18.0	20.9	25.1	932.0	3274.0	5245.0
vie	Vietnamese	Latn	2499.0	943.4	17.1	19.5	22.7	1300.0	3785.0	6283.0
xho	Xhosa	Latn	2498.0	1634.7	19.7	23.6	27.2	917.0	3711.0	5946.0
ydd	Eastern Yiddish	Hebr	2217.0	1249.8	17.4	21.9	27.0	1080.0	4664.0	7094.0
zsm	Standard Malay	Latn	2500.0	1370.4	19.6	23.0	26.5	1232.0	4030.0	6616.0

Table 9: (cont.) Full overview of FiNERWEB.