

# Multilingual and Cross-Lingual Citation Needed Detection on Wikipedia for Lower-Resource Languages

Gerrit Quaremba<sup>1</sup>, Amy Rechkemmer<sup>1</sup>, Elizabeth Black<sup>1</sup>, Denny Vrandečić<sup>2</sup>, Elena Simperl<sup>1</sup>

<sup>1</sup>King’s College London, <sup>2</sup>Wikimedia Foundation

{gerrit.quaremba, amy.rechkemmer, elizabeth.black, elena.simperl}@kcl.ac.uk  
denny@wikimedia.org

## Abstract

In automated fact-checking (AFC), check-worthiness detection identifies claims requiring verification based on domain-specific criteria. On Wikipedia, this task instantiates as Citation Needed Detection (CND), which flags claims lacking supporting citations. However, existing research has largely overlooked lower-resource languages, and recent AFC pipelines rely on large language models (LLMs), which are inaccessible to low-resource organizations. We introduce MCN, a multilingual CND corpus spanning 18 languages across three resource levels, on which we conduct an extensive study of small decoder-based language models (SLMs). Our experiments show that SLMs fine-tuned with an encoder-style objective substantially outperform prompted LLMs across languages. We further present one of the first studies on cross-lingual CND, demonstrating that SLMs fine-tuned solely on English claims surpass LLMs, even with little to no target-language adaptation. Our findings have important implications for lower-resource Wikipedia communities and suggest that compact, task-specific models are preferable to LLMs for CND. We release all data and code at <https://github.com/gerritq/mcn>

## 1 Introduction

The surge of mis- and disinformation has increased the demand for fact-checking (Vlachos and Riedel, 2014). As manual fact-checking is laborious (Nakov et al., 2021), automated fact-checking (AFC) has emerged as a critical approach to support fact-checkers. A crucial early-stage task in AFC is claim check-worthiness detection (CWD), identifying claims that require verification based on domain-specific criteria (Guo et al., 2022).

On Wikipedia, CWD operationalizes as citation needed detection (CND) (Redi et al., 2019; Guo et al., 2022). We present examples of such claims in Figure 2 and Appendix A. Wikipedia editors collectively judge and determine the check-worthiness

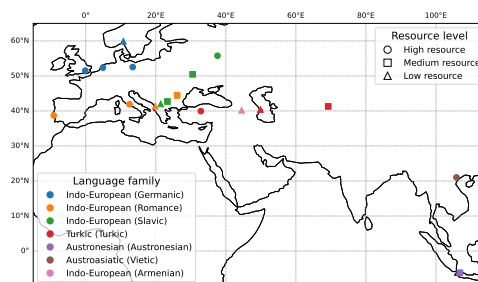


Figure 1: MCN Language Coverage.

of claims by applying Wikipedia’s verifiability policy.<sup>1</sup> When editors encounter claims that violate the policy, they tag them with the CITATION NEEDED template to signal the need for an authoritative citation. This maintenance work is critical for preserving Wikipedia’s knowledge integrity, which has made Wikipedia an essential resource for the artificial intelligence community (Johnson et al., 2024). Yet, CND remains a largely manual task, motivating research that frames it as a supervised binary classification problem (Redi et al., 2019; Halitaj and Zubiaga, 2024; Baigutanova et al., 2026).

However, existing work has primarily studied CND in high-resource languages.<sup>2</sup> This is a critical limitation, as lower-resource language communities typically have fewer active editors and therefore face a disproportionate burden in tagging claims that require citations (Lewoniewski et al., 2017). Moreover, recent AFC pipelines have often relied on large language models (LLMs) for claim detection (e.g., Li et al., 2024a; Wang et al., 2024). Their proprietary nature and high inference costs, however, make LLMs inaccessible to low-resource organizations (Warren et al., 2025).

In this work, we address these gaps by in-

<sup>1</sup>The verifiability policy states that *all claims that are challenged or likely to be challenged must be supported by reliable sources*. We treat each sentence as a claim.

<sup>2</sup>We adopt the language resource taxonomy of Joshi et al. (2020).

roducing **Multilingual Citation Needed (MCN)**, a large-scale CND dataset derived from high-quality Wikipedia articles covering 18 linguistically diverse languages across multiple resource levels (Joshi et al., 2020) (Figure 1). MCN advances prior work by extending coverage to lower-resource languages and by substantially increasing scale (Table 1). On these data, we study small decoder-based language models (SLMs)<sup>3</sup> as cost-effective, open-source alternatives to LLMs for CND. To this end, we specialise SLMs for CND through various supervised fine-tuning (SFT) (Wei et al., 2021) variants with QLoRA (Detmeters et al., 2023), to meet low-resource constraints.

We structure our study around three research questions for multilingual CND. First, initial work demonstrates strong performance of SLMs for English CWD (Majer and Šnajder, 2024; Li et al., 2024b; Bell, 2025). However, their effectiveness in lower-resource languages remains unexplored, and existing work relies on standard SFT. This motivates: **(RQ1)** *How do different SFT variants affect SLM performance across languages?* Second, recent AFC work relies predominantly on LLMs (Dmonte et al., 2024), despite their prohibitive cost for low-resource organizations. We assess whether smaller models offer a viable alternative: **(RQ2)** *How do fine-tuned SLMs compare to prompted LLMs and encoder-based PLMs across languages?* Third, to our knowledge, no prior work has evaluated cross-lingual CND (X-CND). This setting is particularly relevant for lower-resource languages, where labelled data is scarce and supervised training is often infeasible. We therefore ask: **(RQ3)** *How well do English-trained models transfer to lower-resource target languages under zero-shot and few-shot settings?*

Our experiments show that encoder-style fine-tuning of SLMs (Suganthan et al., 2025) yields gains of up to 8 percentage points over generative SFT variants (RQ1). Across all resource levels, fine-tuned SLMs and PLMs substantially outperform prompted LLMs, with PLMs remaining a competitive baseline (RQ2). For cross-lingual transfer (RQ3), English-trained SLMs achieve strong zero-shot performance on parallel data but accuracy drops by about 20 percentage points un-

<sup>3</sup>LM taxonomy is not mutually exclusive (Minaee et al., 2025). We refer to encoder-based pretrained language models (e.g., BERT-based models) as *PLMs*, to decoder-based language models with fewer than 10B parameters as *SLMs* (Belcak et al., 2025), and to larger decoder-based models as *LLMs*.

### ! Citation Needed

*London's gross regional product in 2019 was £503 billion, around a quarter of UK GDP [198].*

### ✓ No Citation Needed

*To explore Pandora, human scientists control Na'vi-human hybrid bodies called "avatars."*

Figure 2: Dataset examples from the Wikipedia articles “London” and “Avatar (2009 film)”. The first sentence requires a citation because it presents a statistical claim, whereas the second does not, as its source (the film itself) is evident to the reader.

der real-world conditions. Notably, encoder-style SLMs fine-tuned solely on English outperform prompted LLMs—even with little to no target-language supervision. We further demonstrate that our findings obtained on high-quality FA articles generalize to randomly sampled Wikipedia claims, underscoring the practical utility of our models.

Our findings have important implications for practitioners aiming to support lower-resource language communities. Crucially, fine-tuning SLMs on high-resource languages alone enables strong cross-lingual transfer, providing a viable path to deployment even when target-language data is scarce. Beyond Wikipedia, our findings reinforce emerging evidence (e.g., Bucher and Martini, 2024; Majer and Šnajder, 2024) that compact, task-specific models are often preferable to general-purpose LLMs.

Our contributions are as follows:

- We introduce **Multilingual Citation Needed (MCN)**, a large-scale CND corpus covering 18 linguistically diverse languages across three resource levels sourced from high-quality Wikipedia articles. MCN extends prior datasets by including lower-resource languages and providing substantially larger samples.
- We study the impact of different SFT variants on adapting SLMs for multilingual CND, and benchmark them against encoder-based PLMs and LLMs. Our results show that fine-tuned models consistently outperform LLM

Dataset	Task	Domain	Languages (resource-level)	Size
NLP4IF 2021 (Shaar et al., 2021)	CD/CWD	Covid-19/Politics	3 (2 high, 1 medium)	1.3K–4K
CitationNeeded (Redi et al., 2019)	CWD	Wikipedia	3 (3 high)	20K
Kazemi et al. (2021)	CD	Covid-19, Politics	5 (2 high, 2 medium, 1 low)	5K
Dutta et al. (2023)	CD	Politics	3 (2 high, 1 medium)	600–1.4K
Halitaj and Zubiaga (2024)	CWD	Wikipedia	3 (2 high, 1 medium)	31K–1.1m
CheckThat 2024 (Hasanain et al., 2024)	CD/CWD	Politics, Covid-19	3 (3 high)	1K–23K
Baigutanova et al. (2026)	CWD	Wikipedia	5 (5 high)	20K
<b>MCN (ours)</b>	CWD	Wikipedia	<b>18 (8 high, 5 medium, 5 low)</b>	<b>2K–1.1m</b>

Table 1: Comparison of multilingual claim detection (CD) and claim check-worthiness detection (CWD) datasets. Resource-level definitions follow Joshi et al. (2020). MCN substantially expands both linguistic and resource-level coverage compared to prior CD/CWD datasets.

baselines.

- To our knowledge, we present the first evaluation of cross-lingual CND. Our results show that SLMs can transfer effectively across languages in zero- and few-shot regimes, highlighting their practical potential for deployment in low-resource settings.

## 2 Related work

**Citation Needed Detection (CND)** CWD is an early-stage task in AFC that aims to identify claims requiring verification based on domain-specific criteria (Guo et al., 2022). Unlike verifiable claim detection (CD), CWD is inherently domain dependent (Konstantinovskiy et al., 2021; Panchendraran and Zubiaga, 2024). Accordingly, prior work has studied CWD across a range of domains, including rumor detection (Zubiaga et al., 2018), politics (Hassan et al., 2015), and Wikipedia (Redi et al., 2019). Redi et al. (2019) formalized the CND task for Wikipedia as a binary supervised classification problem that identifies English claims requiring citations. Halitaj and Zubiaga (2024) extended CND to three smaller non-English Wikipedias, and subsequent work proposed a sample-efficient CND framework (Halitaj and Zubiaga, 2025).<sup>4</sup> Recently, the Wikimedia Foundation expanded CND to several high-resource languages, including Spanish and Italian (Baigutanova et al., 2026). However, existing work largely focuses on high-resource languages, leaving lower-resource language communities underexplored.

**Multilingual Claim Detection** The majority of claim detection research has focused on English (e.g. Arslan et al., 2020; Gangi Reddy et al., 2022). The most prominent non-English claim detection efforts originate from the NLP4IF (Shaar

et al., 2021) and CheckThat! (Hasanain et al., 2024) shared task series. Although a small number of additional multilingual efforts exist (Table 1), several studies have emphasized the need for multilingual claim detection resources (Panchendraran and Zubiaga, 2024; Dmonte et al., 2024; Guo et al., 2022). Regarding cross-lingual approaches, only Panda and Levitan (2021) investigate cross-lingual transfer on the NLP4IF dataset (Shaar et al., 2021) and show that mBERT generalizes well to high-resource languages. Our work addresses the critical lack of multilingual CWD data and enables both multilingual and cross-lingual experiments across languages with varying resource levels.

## 3 Data

We introduce **Multilingual Citation Needed (MCN)**, a CND corpus that spans 18 linguistically diverse languages across three resource levels (Joshi et al., 2020), sampled from high-quality Wikipedia articles. MCN advances prior datasets by incorporating medium- and low-resource languages and by containing a substantially larger number of claims (see Table 1).

**Language Selection** We select 18 languages that span linguistic diversity and represent high-, medium-, and low-resource levels (Joshi et al., 2020) (Figure 1). Our selection is constrained by the availability of sufficiently high-quality Wikipedia articles from which we sample claims.

To ensure linguistic diversity, we sample languages from multiple families, groups, and scripts. We include several Indo-European (IE) language groups (e.g., Germanic, Slavic, and Albanian), covering three scripts (Armenian, Latin, and Cyrillic). To increase family-level diversity, we incorporate languages from the Austronesian, Austroasiatic, and Turkic families. Balancing linguistic diversity with resource availability, we select at least three languages per resource level from related language

<sup>4</sup>Following (Joshi et al., 2020), only Albanian qualifies as a low-resource language in their work.

Language Statistics					No Citation Needed (0)			Citation Needed (1)			Unique Topics	Topic Similarity
Resource	N Featured Articles	Language Group (Family)	Script	N Claims	Avg Len	% Numeric	N Claims	Avg Len	% Numeric			
English (en)	high	6,827	Germanic (IE)	Latin	20,824	20.1	16.8	1,142,036	23.4	39.3	51	0.71
Portuguese (pt)	high	3,485	Romance (IE)	Latin	7,357	22.0	31.6	499,035	25.6	34.9	48	0.94
German (de)	high	2,945	Germanic (IE)	Latin	273,180	18.4	32.6	615,757	19.1	35.2	51	0.96
Russian (ru)	high	2,094	Slavic (IE)	Cyrillic	26,195	17.3	28.5	709,255	19.8	31.4	49	0.92
Italian (it)	high	1,166	Romance (IE)	Latin	17,030	27.6	28.5	180,235	32.2	38.3	36	0.95
Vietnamese (vi)	high	1,044	Mon-Khmer (AA)	Latin	10,664	28.8	25.6	183,366	32.5	40.3	37	0.93
Turkish (tr)	high	573	Oghuz (TRK)	Latin	4,021	14.8	19.7	64,679	16.9	37.2	28	0.78
Dutch (nl)	high	380	Germanic (IE)	Latin	42,863	18.7	21.7	69,490	19.1	26.0	38	0.98
Ukrainian (uk)	mid	1,343	Slavic (IE)	Cyrillic	63,100	17.5	33.5	174,986	18.7	39.1	45	0.96
Romanian (ro)	mid	694	Romance (IE)	Latin	13,104	21.8	26.8	121,558	24.7	36.9	35	0.88
Indonesian (id)	mid	650	Malayo-Polynesian (AN)	Latin	4,446	19.2	30.1	106,202	20.8	34.4	32	0.88
Bulgarian (bg)	mid	190	Slavic (IE)	Cyrillic	12,091	19.0	31.3	27,722	19.9	34.1	27	0.96
Uzbek (uz)	mid	129	Karluk (TRK)	Latin	1,867	14.1	35.1	7,495	16.1	38.5	12	0.65
Norwegian (no)	low	1,406	Germanic (IE)	Latin	96,118	18.2	31.0	192,893	18.6	34.6	51	0.95
Azerbaijani (az)	low	768	Oghuz (TRK)	Latin	23,911	15.5	24.6	146,637	16.1	31.1	35	0.92
Macedonian (mk)	low	395	Slavic (IE)	Cyrillic	33,775	19.9	27.4	60,799	21.6	33.8	37	0.95
Armenian (hy)	low	118	Armenian (IE)	Armenian	6,873	16.5	34.1	25,090	17.6	38.9	27	0.84
Albanian (sq)	low	76	Albanian (IE)	Latin	3,990	21.8	24.8	10,664	24.6	46.0	18	0.38
Total/Average					661,409	19.5	28.0	4,337,899	21.5	36.1	76	0.86

Table 2: Dataset statistics. MCN covers 18 languages spanning multiple resource levels (Joshi et al., 2020). Claims are sourced from Wikipedia featured articles and cover diverse language families and scripts.

groups, including Germanic, Common Turkic, and Slavic.

**Data Collection** For each Wikipedia language edition, we collect claims from *featured articles* (FAs), which editors designate as the highest-quality content on Wikipedia.<sup>5</sup> One FA criterion states that "claims are verifiable against high-quality reliable sources and are supported by inline citations where appropriate."<sup>6</sup> Prior analyses of Wikipedia’s citation quality confirm that FAs are consistently well sourced (Baigutanova et al., 2023). Consequently, we can derive binary CND labels, whether a claim requires a citation, by leveraging the presence or absence of inline citations. Following prior CND work (Redi et al., 2019; Halitaj and Zubiaga, 2024), we label a claim as 1 (Citation Needed) if it contains an inline citation or appears in a paragraph with a citation at the end; we label claims without inline citations as 0 (No Citation Needed). We further augment each claim with contextual information (the preceding and following sentences, when available) and the article’s topic (e.g., Health), which we obtain using Wikipedia’s ORES topic model (Halfaker and Geiger, 2020). We remove all citation markers from claims and apply extensive cleaning to ensure that no citation artifacts remain. Appendix A provides additional details on the data collection and cleaning procedures.

**Dataset Statistics** Table 2 presents dataset statistics. As expected, lower-resource languages contain fewer FAs. However, even among languages

within the same resource level, the number of available articles varies substantially. Across languages, claims that require citations tend to be longer and are more likely to include numeric expressions. These claims originate from articles that cover, on average, 76 unique topics (e.g., Health, Geography). Topic Similarity, computed as the cosine similarity between the topic distributions of claims with and without citations, remains relatively high. This pattern holds across languages except for Albanian, likely due to its considerably smaller number of available FAs. Appendix A provides additional descriptive statistics and example claims.

**Data Quality** A potential concern is that parsing errors during data collection or incorrect citation retrieval may introduce bias. To address this issue, we conduct an extensive manual evaluation of data quality, which we report in Appendix A.

**Data Split** For each language, we split the data at the article level to ensure that test claims originate from articles unseen during training. Across all experiments, we use the same train/dev/test splits of 5k/500/500 claims, balanced by label. We report accuracy as the primary evaluation metric and provide F1 scores in Appendix C.

## 4 Experimental Setup

We first describe the different SFT variants for adapting SLMs to CND, followed by a description of our experimental setup.

### 4.1 SFT Variants

We specialize SLMs via SFT (Wei et al., 2021) using QLoRA (Detmeters et al., 2023), making our experiments reproducible on a single A100 40GB.

<sup>5</sup>[https://en.wikipedia.org/wiki/Category:Featured\\_articles](https://en.wikipedia.org/wiki/Category:Featured_articles)

<sup>6</sup>[https://en.wikipedia.org/wiki/Wikipedia:Featured\\_article\\_criteria](https://en.wikipedia.org/wiki/Wikipedia:Featured_article_criteria)

We explore three SFT variants, with generative and discriminative objectives:

**SFT Variant 1: Full-token loss (FTL)** The FTL is the standard loss in most work, and the default in Huggingface’s SFTTrainer (von Werra et al., 2020). Training minimizes the autoregressive next-token loss over *all* tokens of the prompt (i.e., system, user, and assistant messages). This strategy has been adapted in prior claim detection work (e.g., Li et al., 2024b; Bell, 2025).

**SFT Variant 2: Target-only loss (TOL)** TOL differs from FTL by masking system and user tokens so that model weights update only on assistant (i.e., target) tokens. When fine-tuning models with QLoRA on the MMLU benchmark, Dettmers et al. (2023) show that TOL outperforms FTL. The motivation for applying TOL to CND is to encourage the model to ‘focus’ its learning capacity on predicting the label rather than reconstructing the input prompt.

**SFT Variant 3: Encoder-style (ES)** Rather than using generative SFT, recent work shows that SLMs can be effectively adapted as encoder-style discriminators (e.g., Bolton et al., 2024; Li et al., 2025). We implement this variant similar to Suganthan et al. (2025) by feeding the last-token hidden state into a feed-forward classifier. We fine-tune this variant by concatenating the claim with its context and optimizing it using standard cross-entropy loss.

**SFT Data** For generative SFT, we need to format each data sample into a chat-style format. Given our dataset  $d = \{(c_i, x_i, y_i)\}_{i=1}^N$  consisting of *claim*  $c_i$ , *context*  $x_i$ , and *label*  $y_i$ , we convert each example into a chat-style prompt with system, user, and assistant messages. The system message contains task instructions; the user message concatenates the claim and its context (if available); and the assistant message contains the target label formatted as JSON (e.g., {"label": 1}) to enforce consistent output formatting.

**Models** We use Llama-8B (Grattafiori et al., 2024), Qwen3-8B (Yang et al., 2025) and Aya-8b (Üstün et al., 2024). We omit smaller sizes which we leave for future work.

## 4.2 Experiments

We design an experiment for each research question. Appendix B presents experimental details for

full reproducibility.

**(1) SFT Variants** To evaluate the effectiveness of different SFT variants for adapting SLMs to CND (RQ1), we fine-tune each model for every combination of SFT variant and language using a hyperparameter grid search. We use a simple instruction prompt across all runs, as preliminary experiments show little benefit from prompt engineering.

**(2) Multilingual CND Benchmark** To evaluate the performance of different LM families for CND on MCN (RQ2), we benchmark SLMs against encoder-based PLMs and commercial LLMs. For PLMs, we fully fine-tune mBERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020) in both base and large variants. For LLMs, we evaluate GPT-5.2 (OpenAI, 2025) and Gemini 2.5 Flash (Comanici et al., 2025), combining zero- and few-shot prompting with both instruction-style and verbose prompts, as LLMs may benefit from more detailed task explanations (Majer and Šnajder, 2024).

**(3) X-CND** The ability of trained models to perform cross-lingual CND remains an open question (RQ3). We investigate this question by training mBERT and Llama with all three SFT variants on English as the *source* language. We then evaluate their zero-shot and few-shot transfer performance on all medium- and low-resource *target* languages in our corpus.

For zero-shot transfer, we consider two evaluation settings: (i) **Parallel Test Data**: Because claims in MCN are not parallel, we construct synthetic parallel test sets by translating the English test set into all target languages included in our study.<sup>7</sup> To address potential translation quality issues (Artetxe et al., 2020), we perform extensive automatic quality checks, which we detail in Appendix A. (ii) **Non-parallel Test Data**: Evaluation on non-parallel test data assesses models’ cross-lingual transfer under both language *and* claim distribution shifts (e.g., topics and semantics). This setting poses a more challenging yet more realistic evaluation scenario, as models deployed in practice must classify claims that diverge from their training data distribution.

For few-shot transfer, we follow the strategy by Lauscher et al. (2020) for the non-parallel data setting: each model trained on English is further fine-tuned on a small number of target-language

<sup>7</sup>We use translation-specialized LLMs via Google Cloud Translation.

Model	SFT	High Resource									Medium Resource					Low Resource						
		en	pt	de	ru	it	vi	tr	nl	Avg	uk	ro	id	bg	uz	Avg	no	az	mk	hy	sq	Avg
Aya	FTL	83.28 (±0.71)	56.47 (±1.95)	54.08 (±0.47)	65.73 (±4.81)	59.99 (±4.18)	68.72 (±5.64)	71.53 (±0.31)	53.56 (±1.91)	<b>64.17</b> (±10.24)	58.48 (±3.20)	65.27 (±4.16)	66.79 (±2.75)	54.11 (±0.41)	50.23 (±0.41)	<b>58.98</b> (±7.09)	56.93 (±3.72)	56.20 (±6.30)	56.41 (±4.46)	51.08 (±3.46)	<b>71.40</b> (±3.65)	58.40 (±4.78)
	TOL	82.03 (±1.67)	67.47 (±0.80)	62.66 (±1.11)	73.81 (±0.46)	69.55 (±1.78)	74.95 (±5.02)	<b>74.40</b> (±2.95)	50.20 (±0.17)	<b>69.38</b> (±9.66)	59.52 (±2.61)	74.60 (±1.45)	56.71 (±11.80)	52.14 (±3.27)	50.00 (±0.00)	<b>58.59</b> (±9.70)	51.80 (±2.46)	59.13 (±7.96)	53.12 (±4.70)	47.84 (±7.20)	<b>68.40</b> (±5.09)	56.06 (±8.00)
	ES	<b>89.45</b> (±0.90)	<b>78.58</b> (±1.29)	<b>64.33</b> (±1.31)	<b>78.02</b> (±0.85)	<b>74.14</b> (±1.80)	<b>80.83</b> (±1.95)	72.13 (±0.23)	<b>62.46</b> (±0.81)	<b>74.99</b> (±8.82)	<b>70.21</b> (±1.80)	<b>77.89</b> (±1.49)	<b>70.81</b> (±1.10)	<b>57.85</b> (±1.17)	<b>59.27</b> (±1.72)	<b>67.20</b> (±8.47)	<b>65.33</b> (±0.99)	<b>69.20</b> (±0.53)	<b>63.92</b> (±0.91)	<b>59.59</b> (±1.45)	<b>70.40</b> (±4.78)	<b>65.69</b> (±4.53)
Llama3	FTL	88.64 (±1.00)	77.84 (±1.52)	60.65 (±1.86)	77.15 (±0.87)	73.94 (±1.24)	77.35 (±0.40)	73.40 (±0.53)	60.52 (±0.58)	<b>73.69</b> (±9.38)	69.14 (±2.23)	75.60 (±0.69)	74.55 (±1.22)	<b>60.86</b> (±5.43)	<b>66.73</b> (±6.03)	<b>69.38</b> (±2.57)	62.07 (±2.92)	67.47 (±0.91)	65.93 (±2.31)	58.59 (±2.31)	<b>71.60</b> (±0.92)	65.13 (±5.00)
	TOL	89.58 (±1.84)	77.24 (±0.46)	61.86 (±1.48)	77.35 (±0.87)	<b>74.97</b> (±1.04)	<b>80.63</b> (±1.22)	<b>73.87</b> (±0.58)	60.59 (±1.07)	74.51 (±9.52)	68.54 (±1.39)	<b>76.14</b> (±1.64)	75.15 (±0.53)	60.05 (±0.50)	60.00 (±4.33)	67.98 (±7.82)	63.80 (±0.92)	66.11 (±0.89)	64.99 (±1.74)	59.46 (±6.68)	67.07 (±4.13)	64.28 (±2.96)
	ES	<b>89.65</b> (±0.23)	<b>78.58</b> (±1.03)	<b>66.00</b> (±0.64)	<b>78.89</b> (±1.16)	74.83 (±2.08)	80.09 (±1.29)	71.47 (±1.30)	<b>61.19</b> (±3.18)	<b>75.09</b> (±8.89)	<b>72.01</b> (±1.01)	73.59 (±8.09)	<b>75.55</b> (±1.22)	53.51 (±4.82)	57.67 (±4.96)	<b>66.46</b> (±10.12)	<b>65.73</b> (±2.89)	<b>71.47</b> (±2.27)	<b>60.79</b> (±1.69)	<b>60.79</b> (±1.89)	<b>69.80</b> (±1.91)	<b>66.85</b> (±4.13)
Qwen3	FTL	88.11 (±2.42)	76.77 (±2.15)	56.85 (±0.98)	74.62 (±0.99)	73.46 (±2.09)	79.02 (±1.41)	73.93 (±0.31)	58.52 (±0.60)	<b>72.66</b> (±10.37)	61.55 (±1.94)	74.73 (±3.44)	<b>73.41</b> (±2.82)	60.45 (±0.31)	58.73 (±1.29)	<b>65.78</b> (±7.65)	62.33 (±4.39)	67.40 (±2.21)	58.82 (±2.53)	57.75 (±2.76)	69.27 (±3.13)	63.11 (±5.10)
	TOL	89.11 (±1.36)	<b>77.31</b> (±0.35)	<b>60.72</b> (±1.57)	<b>77.42</b> (±0.83)	73.53 (±2.65)	78.09 (±1.10)	<b>75.53</b> (±0.31)	60.05 (±2.03)	<b>73.97</b> (±9.57)	67.05 (±0.74)	76.48 (±1.31)	<b>71.88</b> (±0.95)	<b>60.99</b> (±1.52)	<b>72.00</b> (±0.80)	<b>69.68</b> (±5.89)	<b>63.67</b> (±1.86)	<b>65.62</b> (±2.22)	60.23 (±1.87)	60.09 (±0.86)	<b>71.00</b> (±1.51)	64.12 (±4.50)
	ES	<b>89.38</b> (±0.53)	76.97 (±0.31)	<b>65.06</b> (±1.14)	<b>77.49</b> (±0.70)	<b>75.58</b> (±0.83)	<b>81.10</b> (±0.12)	71.60 (±1.91)	<b>62.26</b> (±2.62)	<b>74.93</b> (±8.68)	<b>71.08</b> (±2.09)	<b>77.15</b> (±0.42)	70.88 (±4.07)	59.39 (±3.31)	58.60 (±4.20)	67.42 (±8.10)	<b>66.67</b> (±2.27)	<b>71.47</b> (±1.97)	<b>63.11</b> (±1.48)	<b>62.26</b> (±1.50)	68.13 (±2.00)	<b>66.33</b> (±3.70)

Table 3: Monolingual CND accuracy of SLMs across SFT variants (FTL=full-token loss, TOL=target-only loss, ES=encoder-style). **Bold** denotes the best score per model–language pair. Parentheses indicate standard deviation. ES SFT consistently outperforms generative SFT, while TOL yields modest gains over FTL in some settings.

examples ( $N < 500$ ). While this approach has proven effective for PLMs across a range of NLP tasks (Lauscher et al., 2020), its effectiveness for CND remains unclear and, to our knowledge, it has not previously been applied to SLMs.

## 5 Results

### Experiment 1: SFT Variants

**Encoder-style tuning mostly outperforms other SFT variants for multilingual CND across language resource levels.** Table 4 reports CND accuracy for SLMs across SFT variants and languages of all resource levels. Across nearly all model-resource combinations, ES tuning achieves higher average accuracy than generative SFT variants. This effect is most pronounced for Aya, where ES tuning outperforms generative variants on almost all individual languages (e.g., by an average of 10 percentage points compared to FTL on high-resource languages). For Qwen and Llama, ES tuning also outperforms other variants on both high- and low-resource languages, although the gains are smaller (e.g., an average improvement of 3 percentage points for Llama on low-resource languages). On medium-resource languages, generative tuning variants achieve higher average performance. However, this result is driven by a single outlier (Uzbek) where generative tuning yields accuracy gains of up to 13 percentage points over ES tuning. For most other medium-resource languages, ES tuning remains competitive with generative SFT. Among the generative variants, TOL training yields slightly higher performance than FTL on high-resource languages across models. On medium- and low-resource languages, TOL and FTL perform comparably, with differences of only 1–2 percentage points.

### Experiment 2: Multilingual CND Benchmark

**SLMs considerably outperform LLMs on CND across most languages.** Table 4 reports accuracy benchmarking results for all LMs in our study. For SLMs, we report results for the ES variant. Across resource levels, SLMs outperform LLMs by an average margin of 8–12%. Although LLMs achieve competitive results for Turkish and Bulgarian, ranking second best and best, SLMs consistently outperform LLMs even on most individual languages across all resource levels. LLM performance degrades as language resources decrease. The best-performing LLM, Gemini 2.5, achieves only 59.44% accuracy on low-resource languages, compared to 66.85% for Qwen. A similar trend appears for SLMs and PLMs when moving from high- to medium-resource languages. However, performance remains relatively stable from medium- to low-resource languages (e.g., Llama achieves 66.46% accuracy on medium-resource languages and 66.85% on low-resource languages). With detection accuracy ranging between 53–70% for medium- and low-resource languages, these results highlight the need for further research focused on low-resource language communities. We note that performance differences across languages stem from both language-specific factors and variations in claim distributions.<sup>8</sup> We are able to analyze these effects in the X-CND experiments presented in the next section. **Encoder-based PLMs remain a competitive baseline for multilingual CND.** Encoder-based PLMs continue to provide strong baselines for multilingual CND. mBERT achieves the highest average accuracy on medium- and low-resource languages, with moderate gains of 1–2 percentage points over SLMs. Compared

<sup>8</sup>A notable outlier is German. We provide a qualitative analysis of misclassified instances in Appendix C.

	High Resource									Medium Resource					Low Resource							
	en	pt	de	ru	it	vi	tr	nl	Avg	uk	ro	id	bg	uz	Avg	no	az	mk	hy	sq	Avg	Avg
<b>Decoder-based LLMs</b>																						
<i>Zero-shot</i>																						
GPT-5 (Instruct)	75.00	58.80	54.40	65.20	60.52	64.40	69.20	50.60	<u>62.27</u> (±1.90)	53.20	65.00	53.60	53.00	56.60	<u>56.28</u> (±3.09)	51.80	59.00	56.00	50.00	53.20	<u>54.00</u> (±3.35)	<u>58.31</u> (±3.98)
GPT-5 (Verbose)	80.60	57.00	51.60	62.40	55.20	66.40	69.40	50.40	<u>61.62</u> (±1.29)	52.20	64.60	55.80	53.20	54.20	<u>56.00</u> (±3.99)	52.80	61.20	53.40	50.60	51.60	<u>53.92</u> (±4.21)	<u>57.92</u> (±8.09)
Gemini 2.5 (Instruct)	75.80	59.20	56.20	65.80	64.20	63.20	62.40	51.20	<u>62.25</u> (±1.20)	55.60	63.40	57.40	57.00	60.20	<u>58.72</u> (±3.10)	54.00	54.60	55.00	50.60	55.60	<u>52.96</u> (±1.97)	<u>58.97</u> (±6.11)
Gemini 2.5 (Verbose)	79.20	62.80	54.40	62.00	52.60	67.40	68.40	52.00	<u>62.35</u> (±3.39)	54.80	62.40	58.40	58.00	55.20	<u>57.76</u> (±3.89)	53.00	58.40	53.60	49.40	51.00	<u>53.08</u> (±3.43)	<u>58.50</u> (±7.59)
<i>Few-shot</i>																						
GPT-5 (Instruct)	82.00	58.80	54.00	68.00	68.61	69.00	<u>72.20</u>	53.31	<u>65.74</u> (±1.79)	53.60	71.40	56.80	57.40	67.00	<u>61.24</u> (±2.37)	52.60	61.20	55.40	52.80	55.00	<u>55.40</u> (±3.48)	<u>61.62</u> (±8.87)
GPT-5 (Verbose)	80.40	58.20	50.80	62.20	57.00	68.74	68.94	51.20	<u>62.19</u> (±1.09)	53.31	65.80	57.40	53.60	56.00	<u>57.22</u> (±5.09)	52.60	61.12	53.40	51.60	50.60	<u>53.86</u> (±4.19)	<u>58.49</u> (±8.09)
Gemini 2.5 (Instruct)	80.60	58.00	59.80	67.54	66.60	68.20	69.20	55.00	<u>65.62</u> (±3.05)	55.60	62.20	62.60	<b>63.00</b>	<u>67.20</u>	<u>62.12</u> (±3.16)	54.60	62.20	57.00	59.00	64.40	<u>59.44</u> (±3.33)	<u>62.93</u> (±8.45)
Gemini 2.5 (Verbose)	81.00	63.00	50.40	61.60	53.60	71.00	67.60	53.00	<u>62.65</u> (±1.40)	55.60	62.40	63.20	55.00	61.40	<u>59.52</u> (±3.91)	53.80	58.80	56.20	53.40	54.40	<u>55.32</u> (±3.22)	<u>59.74</u> (±7.68)
<b>Encoder-based PLMs</b>																						
XLM-R-base	87.91 (±0.31)	73.49 (±2.41)	63.79 (±1.68)	77.02 (±1.14)	68.59 (±2.88)	73.15 (±1.45)	<b>72.87</b> (±0.58)	55.11 (±5.22)	<u>71.49</u> (±3.60)	65.06 (±2.14)	71.84 (±5.25)	64.80 (±3.41)	53.57 (±3.07)	50.00 (±0.00)	<u>61.05</u> (±9.01)	63.07 (±0.99)	62.73 (±0.85)	62.84 (±1.14)	<u>62.93</u> (±0.77)	<u>66.60</u> (±7.35)	<u>63.63</u> (±1.66)	<u>66.48</u> (±9.21)
XLM-R-large	<b>89.65</b> (±0.23)	<u>77.04</u> (±2.54)	<u>61.06</u> (±1.57)	<u>77.62</u> (±0.94)	<u>73.18</u> (±0.94)	<u>76.82</u> (±0.83)	<u>71.40</u> (±1.91)	54.24 (±7.35)	<u>72.63</u> (±0.84)	<u>62.86</u> (±2.84)	<b>79.50</b> (±1.59)	<u>68.14</u> (±3.22)	55.98 (±0.85)	50.00 (±0.00)	<u>63.30</u> (±11.38)	<u>66.80</u> (±1.49)	64.80 (±1.11)	63.85 (±1.31)	<u>62.32</u> (±1.45)	<u>66.33</u> (±1.82)	<u>64.82</u> (±1.38)	<u>67.86</u> (±10.18)
mBert	88.71 (±0.61)	<b>78.85</b> (±0.61)	64.73 (±1.64)	76.55 (±0.72)	73.39 (±1.13)	78.02 (±2.45)	71.53 (±1.27)	<b>63.33</b> (±1.31)	<u>74.39</u> (±8.17)	70.61 (±2.82)	<u>78.29</u> (±1.57)	<u>73.15</u> (±0.87)	<u>59.65</u> (±1.30)	<b>68.13</b> (±14.20)	<u>69.97</u> (±2.21)	<u>69.27</u> (±0.42)	71.30 (±1.48)	63.31 (±2.26)	<b>63.53</b> (±1.50)	<u>70.47</u> (±1.50)	<u>67.57</u> (±3.76)	<u>71.27</u> (±7.72)
<b>Decoder-based SLMs</b>																						
Aya (ES)	<u>89.45</u> (±0.30)	<u>78.58</u> (±1.29)	64.33 (±1.31)	<u>78.02</u> (±0.83)	74.14 (±1.80)	<u>80.83</u> (±1.95)	72.13 (±0.23)	<u>62.46</u> (±0.81)	<u>74.99</u> (±8.22)	70.21 (±1.55)	77.89 (±1.49)	70.81 (±1.10)	57.85 (±1.17)	59.27 (±1.72)	<u>67.20</u> (±8.47)	65.33 (±0.99)	69.20 (±0.53)	63.92 (±0.91)	59.59 (±1.45)	<u>70.40</u> (±3.78)	<u>65.69</u> (±4.33)	<u>70.24</u> (±8.48)
Llama3 (ES)	<b>89.65</b> (±1.03)	<u>78.58</u> (±1.63)	<b>66.00</b> (±1.64)	<b>78.89</b> (±1.10)	<u>74.83</u> (±2.08)	80.09 (±1.29)	71.47 (±1.30)	61.19 (±1.18)	<b>75.09</b> (±8.88)	<b>72.01</b> (±1.91)	73.59 (±0.89)	<b>75.55</b> (±1.91)	53.51 (±0.90)	57.67 (±1.96)	<u>66.46</u> (±10.38)	<u>65.73</u> (±2.27)	<b>71.47</b> (±1.49)	<b>66.47</b> (±1.39)	60.79 (±1.91)	69.80 (±1.91)	<u>66.85</u> (±3.76)	<u>70.72</u> (±8.21)
Qwen3 (ES)	89.38 (±0.53)	76.97 (±0.31)	<u>65.06</u> (±1.14)	77.49 (±0.70)	<b>75.58</b> (±0.93)	<b>81.10</b> (±0.12)	71.60 (±1.91)	62.26 (±2.62)	<u>74.93</u> (±8.68)	<u>71.08</u> (±2.09)	77.15 (±0.42)	70.88 (±4.07)	59.39 (±3.31)	58.60 (±4.20)	<u>67.42</u> (±8.10)	<u>66.67</u> (±2.27)	<b>71.47</b> (±1.67)	63.11 (±1.48)	62.26 (±1.50)	68.13 (±2.00)	<u>66.33</u> (±3.76)	<u>70.45</u> (±8.21)

Table 4: Monolingual CND accuracy across LM families (ES=encoder-style; Verbose and Instruct denote prompt types). **Bold** and underlined scores mark the best and second-best results, respectively. Parentheses indicate standard deviation. Fine-tuned encoder-based PLMs and decoder-based SLMs substantially outperform LLMs across language resource levels.

to XLM-RoBERTa, including its larger variants, SLMs achieve higher accuracy by 1–4 percentage points.<sup>9</sup> mBERT’s strong performance may stem from its exclusive pretraining on multilingual Wikipedia data, which provides near-perfect domain alignment.

### Experiment 3: X-CND




Figure 3: Zero-shot CND accuracy by test data setting, averaged across language resource groups. SLMs across SFT variants exhibit strong zero-shot capabilities.

**Fine-tuned models demonstrate strong zero-shot capabilities, but performance drops by up to 20 percentage points under real-world conditions. SLMs outperform PLMs by about 10 percentage points.** Figure 3 reports accuracy averaged across test sets and language resource levels. On parallel data, fine-tuned models exhibit strong zero-shot generalization: SLMs achieve up to 87% accuracy, while mBERT reaches approximately 74% on

<sup>9</sup>This difference is partly attributable to XLM-RoBERTa failing on Uzbek. Despite extensive exploration of random seeds and hyperparameter configurations, the model continued to fail on this language.

medium-resource languages, with only marginal declines for low-resource languages. However, on non-parallel data, accuracy drops by 15–20 percentage points across all models. Across all settings, SLMs outperform mBERT by 5–10 percentage points and exhibit lower variance, as indicated by the error bars. Notably, performance differences across SFT variants remain small (within 2 percentage points), suggesting that the choice of SFT variant matters less for zero-shot transfer than for monolingual CND. **QLoRA-fine-tuned SLMs require slower learning to achieve few-shot gains.** For few-shot transfer, we continue fine-tuning the English-trained models on small amounts of target-language data, following Lauscher et al. (2020). However, we find that few-shot adaptation does not reliably improve SLM performance when using the higher learning rates commonly adopted in QLoRA fine-tuning (Dettmers et al., 2023). We therefore treat the learning rate as a tunable hyperparameter during this second-stage adaptation. Appendix C reports the full results and shows that substantially lower learning rates are necessary to obtain consistent few-shot gains for QLoRA-fine-tuned SLMs. **Encoder-style SFT yields the most reliable few-shot gains among SLMs, while mBERT benefits the most from few-shot tuning overall.** Table 5 reports few-shot accuracy results on non-parallel data. For each model, we report results obtained with the optimal learning rate, as described above. Among the SFT variants, only ES tuning exhibits consistent and substantial gains from few-shot adaptation.

Model	Shots	Medium-Resource					$\Delta$ Avg	Low-Resource					$\Delta$ Avg
		uk	ro	id	bg	uz		no	az	mk	hy	sq	
Llama3 (FTL)	0	60.45 ( $\pm 0.70$ )	74.38 ( $\pm 2.02$ )	66.87 ( $\pm 1.95$ )	62.53 ( $\pm 1.74$ )	71.13 ( $\pm 3.00$ )		58.40 ( $\pm 1.78$ )	64.86 ( $\pm 1.86$ )	58.75 ( $\pm 1.21$ )	62.88 ( $\pm 4.22$ )	68.07 ( $\pm 2.57$ )	
	$\Delta$ 50	+1.14 ( $\pm 0.83$ )	-0.54 ( $\pm 0.94$ )	+0.87 ( $\pm 1.75$ )	-0.47 ( $\pm 2.05$ )	-0.33 ( $\pm 2.60$ )	+0.13	-0.40 ( $\pm 1.39$ )	+0.07 ( $\pm 1.51$ )	+1.88 ( $\pm 0.91$ )	0.00 ( $\pm 2.73$ )	-0.40 ( $\pm 3.41$ )	+0.23
	$\Delta$ 100	-0.47 ( $\pm 1.73$ )	-0.67 ( $\pm 0.30$ )	+2.00 ( $\pm 1.36$ )	+1.20 ( $\pm 0.80$ )	+0.60 ( $\pm 2.02$ )	+0.53	+0.60 ( $\pm 1.60$ )	-0.80 ( $\pm 1.10$ )	+1.61 ( $\pm 2.13$ )	-0.76 ( $\pm 3.03$ )	-1.13 ( $\pm 2.50$ )	-0.10
	$\Delta$ 250	-0.94 ( $\pm 1.06$ )	-0.14 ( $\pm 1.34$ )	+3.07 ( $\pm 2.09$ )	+0.47 ( $\pm 0.46$ )	-0.20 ( $\pm 1.29$ )	+0.45	+0.33 ( $\pm 0.46$ )	-0.47 ( $\pm 1.50$ )	-1.14 ( $\pm 1.83$ )	0.00 ( $\pm 2.27$ )	+0.20 ( $\pm 2.21$ )	-0.21
	$\Delta$ 500	+0.13 ( $\pm 2.84$ )	-0.66 ( $\pm 0.42$ )	+2.61 ( $\pm 1.63$ )	+0.94 ( $\pm 0.50$ )	+0.51 ( $\pm 2.62$ )	+0.71	+0.27 ( $\pm 1.10$ )	-0.20 ( $\pm 2.25$ )	+0.94 ( $\pm 1.21$ )	+5.56 ( $\pm 1.91$ )	+0.33 ( $\pm 2.46$ )	+1.38
Llama3 (TOL)	0	57.78 ( $\pm 1.97$ )	73.79 ( $\pm 0.73$ )	68.14 ( $\pm 1.97$ )	62.73 ( $\pm 1.84$ )	69.80 ( $\pm 1.11$ )		58.07 ( $\pm 1.03$ )	64.89 ( $\pm 1.67$ )	57.95 ( $\pm 2.10$ )	56.30 ( $\pm 5.54$ )	66.07 ( $\pm 0.90$ )	
	$\Delta$ 50	+0.63 ( $\pm 0.99$ )	+0.20 ( $\pm 0.29$ )	-1.10 ( $\pm 0.99$ )	+0.20 ( $\pm 0.57$ )	+0.80 ( $\pm 1.41$ )	+0.15	-0.27 ( $\pm 2.51$ )	-0.47 ( $\pm 1.93$ )	+0.94 ( $\pm 0.76$ )	+0.05 ( $\pm 5.05$ )	-0.07 ( $\pm 3.22$ )	+0.04
	$\Delta$ 100	+0.73 ( $\pm 1.42$ )	+0.10 ( $\pm 2.42$ )	+0.30 ( $\pm 0.71$ )	+0.60 ( $\pm 0.28$ )	+1.20 ( $\pm 1.13$ )	+0.59	-0.47 ( $\pm 2.03$ )	-0.73 ( $\pm 2.64$ )	+1.01 ( $\pm 2.44$ )	+0.29 ( $\pm 4.78$ )	+0.73 ( $\pm 3.83$ )	+0.17
	$\Delta$ 250	+1.54 ( $\pm 0.57$ )	+1.11 ( $\pm 0.14$ )	-1.10 ( $\pm 1.84$ )	+1.20 ( $\pm 1.98$ )	-1.80 ( $\pm 1.13$ )	+0.19	-0.47 ( $\pm 1.97$ )	-1.56 ( $\pm 1.71$ )	+0.60 ( $\pm 1.01$ )	+2.94 ( $\pm 2.14$ )	+1.27 ( $\pm 3.21$ )	+0.56
	$\Delta$ 500	+1.14 ( $\pm 0.28$ )	+1.11 ( $\pm 0.71$ )	+0.20 ( $\pm 3.12$ )	-0.50 ( $\pm 2.41$ )	+1.30 ( $\pm 2.97$ )	+0.65	+0.20 ( $\pm 2.14$ )	-1.16 ( $\pm 1.51$ )	+0.67 ( $\pm 1.18$ )	+4.53 ( $\pm 2.67$ )	+1.67 ( $\pm 1.92$ )	+1.18
Llama3 (ES)	0	64.93 ( $\pm 1.51$ )	75.20 ( $\pm 1.40$ )	71.68 ( $\pm 1.75$ )	64.40 ( $\pm 3.64$ )	64.73 ( $\pm 1.62$ )		60.80 ( $\pm 1.11$ )	65.73 ( $\pm 2.23$ )	61.90 ( $\pm 0.81$ )	53.71 ( $\pm 1.75$ )	69.33 ( $\pm 1.70$ )	
	$\Delta$ 50	+0.53 ( $\pm 1.21$ )	+0.60 ( $\pm 0.92$ )	-0.40 ( $\pm 0.70$ )	+2.54 ( $\pm 1.64$ )	+2.60 ( $\pm 5.25$ )	+1.18	+0.20 ( $\pm 1.11$ )	-0.67 ( $\pm 1.33$ )	-1.07 ( $\pm 2.07$ )	+5.88 ( $\pm 1.45$ )	+0.73 ( $\pm 1.86$ )	+1.01
	$\Delta$ 100	-0.20 ( $\pm 1.22$ )	-0.13 ( $\pm 0.51$ )	+1.27 ( $\pm 1.64$ )	+2.07 ( $\pm 2.55$ )	+5.07 ( $\pm 3.94$ )	+1.61	-0.13 ( $\pm 1.10$ )	-0.40 ( $\pm 1.22$ )	-0.34 ( $\pm 1.41$ )	+3.54 ( $\pm 0.61$ )	+0.73 ( $\pm 1.75$ )	+0.68
	$\Delta$ 250	+0.53 ( $\pm 1.17$ )	+0.13 ( $\pm 0.62$ )	+2.61 ( $\pm 1.33$ )	+1.67 ( $\pm 2.21$ )	+5.33 ( $\pm 4.24$ )	+2.06	+1.27 ( $\pm 1.96$ )	+0.40 ( $\pm 2.02$ )	-2.08 ( $\pm 1.90$ )	+4.48 ( $\pm 1.67$ )	+1.80 ( $\pm 0.95$ )	+1.17
	$\Delta$ 500	+3.21 ( $\pm 0.80$ )	-0.54 ( $\pm 0.93$ )	+2.94 ( $\pm 1.81$ )	+1.40 ( $\pm 1.82$ )	+7.73 ( $\pm 3.60$ )	+2.95	+2.33 ( $\pm 1.15$ )	+2.07 ( $\pm 1.60$ )	-1.14 ( $\pm 1.92$ )	+4.81 ( $\pm 1.40$ )	+2.60 ( $\pm 1.67$ )	+2.13
mBert	0	55.04 ( $\pm 1.33$ )	65.79 ( $\pm 6.87$ )	62.66 ( $\pm 4.32$ )	54.98 ( $\pm 2.32$ )	63.93 ( $\pm 5.91$ )		54.00 ( $\pm 2.16$ )	58.93 ( $\pm 4.35$ )	53.66 ( $\pm 2.82$ )	56.65 ( $\pm 3.33$ )	61.33 ( $\pm 5.83$ )	
	$\Delta$ 50	-0.73 ( $\pm 0.53$ )	+6.05 ( $\pm 1.16$ )	-2.20 ( $\pm 5.47$ )	+2.94 ( $\pm 4.74$ )	+2.47 ( $\pm 6.30$ )	+1.70	-1.60 ( $\pm 3.86$ )	+3.27 ( $\pm 0.60$ )	+2.28 ( $\pm 3.43$ )	-1.60 ( $\pm 2.06$ )	+7.93 ( $\pm 3.63$ )	+2.06
	$\Delta$ 100	+0.01 ( $\pm 0.61$ )	+6.59 ( $\pm 5.09$ )	+3.87 ( $\pm 5.95$ )	+2.47 ( $\pm 3.25$ )	+0.60 ( $\pm 4.98$ )	+2.85	-0.40 ( $\pm 2.91$ )	+1.33 ( $\pm 2.53$ )	+3.49 ( $\pm 1.22$ )	+2.07 ( $\pm 4.67$ )	+10.27 ( $\pm 3.86$ )	+3.35
	$\Delta$ 250	+1.67 ( $\pm 1.31$ )	+8.80 ( $\pm 1.12$ )	+3.81 ( $\pm 5.62$ )	+4.28 ( $\pm 2.69$ )	+2.87 ( $\pm 7.30$ )	+4.28	+0.33 ( $\pm 0.42$ )	+3.27 ( $\pm 1.74$ )	+1.88 ( $\pm 1.26$ )	+4.74 ( $\pm 2.45$ )	+4.00 ( $\pm 5.35$ )	+2.84
	$\Delta$ 500	+3.27 ( $\pm 4.34$ )	+3.76 ( $\pm 4.74$ )	-0.33 ( $\pm 2.84$ )	+3.34 ( $\pm 2.11$ )	+6.27 ( $\pm 8.62$ )	+3.26	+3.00 ( $\pm 1.06$ )	+4.00 ( $\pm 1.86$ )	+1.88 ( $\pm 1.05$ )	+5.88 ( $\pm 3.11$ )	+11.47 ( $\pm 1.93$ )	+5.24

Table 5: Few-shot CND accuracy on non-parallel data by model and number of shots.  $\Delta$  denotes gains over zero-shot accuracy. Parentheses report standard deviations. All models are fine-tuned on English claims only. Only Llama (ES) and mBERT benefit from few-shot fine-tuning. Notably, Llama (ES) outperforms LLMs (Table 4) even with no or minimal target-language supervision.

In contrast, continued fine-tuning with FTL and TOL proves largely ineffective on average, yielding only marginal improvements (less than 1% accuracy). By comparison, mBERT consistently benefits from few-shot tuning, particularly on lower-resource languages (e.g., a gain of 11.47 percentage points on Albanian with only 500 samples). Overall, few-shot tuning improves mBERT more than it improves ES-tuned SLMs.

**English-only ES-fine-tuned SLMs surpass LLMs, even with zero or minimal target-language supervision.** Compared to the monolingual evaluation of Experiment 2, ES-tuned Llama outperforms prompted LLMs on nearly all target languages, even under zero-shot or minimal target-language supervision. For example, on Uzbek, Llama (ES) achieves 64.73% accuracy in the zero-shot setting and 72.46% with  $k=500$  target-language examples, compared to 67.20% for the best-performing LLM (Gemini 2.5, few-shot). This result holds for all languages except Armenian. Notably, Llama is not only considerably smaller but also fine-tuned exclusively on English claims for this cross-lingual task.

### Are results generalizable to non-FA articles?

A potential concern is that models trained on claims from FAs may not generalize to non-FA articles, for example due to less standardized editorial styles. To assess this, we collected claims from randomly sampled Wikipedia articles (RA), selecting two languages per resource level, and constructed balanced test sets ( $n = 100$  per language). Because non-FA articles are often under-referenced, leading

	en	it	uk	ro	no	az
RA	62.63	61.00	71.43	64.65	70.00	60.61
RA (corrected labels)	86.87	70.00	72.45	68.68	68.00	64.65
FA baseline	89.65	74.83	72.01	73.59	65.73	71.47
RA Distribution (0/1)	14/85	11/89	6/92	7/92	8/92	2/97

Table 6: Monolingual CND accuracy on the Random Articles (RA) test set ( $n = 100$  per language) using Llama 3 (ES). Due to under-referencing in random articles, we manually corrected the labels.

to a high number of false negatives, we manually corrected all 600 labels.

Table 6 reports results for (i) the originally sampled claims, (ii) the label-corrected claims, and (iii) the FA baseline from Table 4. We find that, for most languages, model performance on RA claims is comparable to, or slightly lower than, performance on FA claims. The discrepancy between results on uncorrected and corrected RA labels is explained by the high number of false negatives in the original annotations. After correction, the label distribution becomes more imbalanced; nevertheless, model performance remains close to FA-level detection. Overall, these results indicate that models trained on FA data generalize well to claims from randomly sampled Wikipedia articles, supporting their practical utility for the Wikipedia community.

## 6 Discussion

**Choosing the right SFT variant for SLMs can yield substantial performance gains for multilingual CND (RQ1).** Experiment 1 shows that ES fine-tuning consistently outperforms generative variants for monolingual CND, although TOL can

benefit certain model–language combinations. Recent work explores adaptations of decoder-based models for encoder-like tasks (Xu et al., 2025; Suganthan et al., 2025) and demonstrates that encoder-style training can be effective (Bolton et al., 2024; Li et al., 2025). Our results align with these findings and show that ES training performs well for CND across languages. For practitioners, ES tuning requires only minimal modifications to the base model while delivering reliable performance gains. In contrast, under generative SFT, the optimal variant depends on the specific model-language combination.

**Supervised models substantially outperform LLMs (RQ2).** Both fine-tuned PLMs and SLMs consistently outperform LLMs on multilingual CND. This result extends prior findings on English CWD (Majer and Šnajder, 2024; Bell, 2025; Li et al., 2024b) to the multilingual setting and aligns with evidence from broader NLP classification tasks (e.g., Bucher and Martini, 2024). Our findings show that supervised models are not only more compact and cost-effective than LLMs but also deliver substantially stronger performance for CND. Their relatively low hardware requirements make them particularly well suited for deployment in resource-constrained environments such as Wikipedia. Nonetheless, the substantial performance variation across resource levels (53-75%) indicates that multilingual CND remains an open and challenging research problem.

**Encoder-style SLMs exhibit strong zero-shot cross-lingual transfer and outperform LLMs with little to no target-language supervision (RQ3).** Across SFT variants, SLMs demonstrate robust zero-shot generalization on parallel data. However, performance drops by up to 20 percentage points under real-world conditions, exposing a critical limitation for practical deployment. This finding highlights the need for future X-CND research to develop more robust strategies that address simultaneous language and claim-distribution shifts.

Crucially, ES-fine-tuned SLMs trained exclusively on English claims outperform prompted LLMs even with no or minimal target-language adaptation. This result has direct implications for real-world deployment: cross-lingual transfer of specialized SLMs trained on high-resource languages emerges as a viable strategy when target-language data is scarce. For practitioners, this approach offers an affordable deployment pathway to support low-resource communities in CND.

## 7 Conclusion

We introduce MCN, a multilingual citation needed detection (CND) dataset that covers Wikipedia claims from 18 linguistically diverse languages across three resource levels. We study decoder-based small language models (SLMs) as alternatives to LLMs for CND, specialized through supervised fine-tuning. We demonstrate that SLMs outperform LLMs on monolingual CND across languages and that encoder-style tuning enables strong zero- and few-shot cross-lingual CND transfer. Our findings have important implications for supporting low-resource language communities in CND.

### Limitations

Our work has several limitations.

**Language selection concentrates on Indo-European languages.** Our selected languages are 70% Indo-European. While we vary language groups and also cover other families, future work can focus on language families such as Sino-Tibetan or Semitic.

**Insights into SLM performance are limited to QLoRA-based SFT.** To ensure practical relevance, we fine-tune SLMs exclusively using QLoRA. As a result, our insights into the effects of SFT strategies and the cross-lingual capabilities of SLMs for CND are limited to this PEFT setting. Performance dynamics may differ under alternative parameter-efficient fine-tuning approaches. Nevertheless, QLoRA is particularly well suited to this study due to its simplicity and efficiency under quantisation. Moreover, we restrict our analysis to 8B-parameter SLMs. Future work should explore a broader range of model sizes and PEFT strategies to assess the generality of our findings.

**Limited generalizability to other domains.** CWD is inherently domain-dependent, as domain-specific criteria determine what constitutes a check-worthy claim. Accordingly, our findings are limited to the Wikipedia domain and may not directly transfer to other common CWD settings, such as political debates. Nevertheless, the strong monolingual and cross-lingual CND performance observed on the challenging Wikipedia domain suggests that comparable gains may be achievable in other application domains.

**Real-world citation practices deviate from labeling heuristics.** We label claims as requiring

a citation or not based on the presence or absence of inline citations, a heuristic following prior CND work (Redi et al., 2019; Halitaj and Zubiaga, 2024; Baigutanova et al., 2026). Nevertheless, in practice, real-world editorial practices deviate from that heuristic. For example, a citation may appear only in the final paragraph of a multi-paragraph passage when it is clear to support all paragraphs. We illustrate this challenge for German Wikipedia in Appendix C. We believe that devising more robust labeling heuristics, for example through the use of soft labels, is an important direction for future work.

## Ethics Statement

**Biases** Although our language selection spans a wide range of language families and writing systems, it is skewed toward Indo-European languages. In addition, Featured Wikipedia articles, while covering diverse topics, tend to emphasize subjects that are of particular interest to the corresponding language communities at a given time, which may introduce topical biases.

**Potential Harms** We are not aware of any direct harms arising from our dataset or experiments. Our work is motivated by the goal of supporting Wikipedia editors, particularly in low-resource language communities, by improving automated citation needed detection.

**Licensing** We collect only publicly available Wikipedia content, which is licensed under CC BY-SA. We release all collected data, including translated text obtained via the Google Translate API, under the same CC BY-SA license to support future research.

**Reproducibility** We provide full details required to reproduce our results in Appendix B, including hyperparameters, hardware requirements, and fine-tuning procedures. All models are trained three times with different random seeds, and we report mean performance and standard deviations to enhance the reliability of our findings.

**AI Assistants** We use LLMs to proofread our writing by prompting it to correct grammar and spelling.

## References

Fatma Arslan, Naemul Hassan, Chengkai Li, and Mark Tremayne. 2020. A benchmark dataset of check-

worthy factual claims. In *Proceedings of the international AAAI conference on web and social media*, volume 14, pages 821–829.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2020. Translation artifacts in cross-lingual transfer learning. *arXiv preprint arXiv:2004.04721*.

Aitolkyn Baigutanova, Jaehyeon Myung, Diego Saez-Trumper, Ai-Jou Chou, Miriam Redi, Changwook Jung, and Meeyoung Cha. 2023. Longitudinal assessment of reference quality on wikipedia. In *Proceedings of the ACM web conference 2023*, pages 2831–2839.

Aitolkyn Baigutanova, Francisco Navas, Pablo Aragón, Mykola Trokhymovych, Muniza Aslam, Ai-Jou Chou, Miriam Redi, and Diego Sáez Trumper. 2026. [Multilingual reference need assessment system for wikipedia](#). In *Proceedings of the ACM Web Conference 2026*, WWW '26, page 8191–8199, New York, NY, USA. Association for Computing Machinery.

Peter Belcak, Greg Heinrich, Shizhe Diao, Yonggan Fu, Xin Dong, Saurav Muralidharan, Yingyan Celine Lin, and Pavlo Molchanov. 2025. [Small language models are the future of agentic ai](#). *Preprint*, arXiv:2506.02153.

Andrew Bell. 2025. Less can be more: An empirical evaluation of small and large language models for sentence-level claim detection. In *Proceedings of the Eighth Fact Extraction and VERification Workshop (FEVER)*, pages 85–90.

Elliot Bolton, Abhinav Venigalla, Michihiro Yasunaga, David Hall, Betty Xiong, Tony Lee, Roxana Daneshjou, Jonathan Frankle, Percy Liang, Michael Carbin, and 1 others. 2024. Biomedlm: A 2.7 b parameter language model trained on biomedical text. *arXiv preprint arXiv:2403.18421*.

Martin Juan José Bucher and Marco Martini. 2024. Fine-tuned ‘small’ llms (still) significantly outperform zero-shot generative ai models in text classification. *arXiv preprint arXiv:2406.08660*.

Eirini Chatzikoumi. 2020. [How to evaluate machine translation: A review of automated and human metrics](#). *Natural Language Engineering*, 26(2):137–161.

Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, and 3416 others. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *Preprint*, arXiv:2507.06261.

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](#). *Preprint*, arXiv:1911.02116.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#). *Preprint*, arXiv:2305.14314.
- 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.
- Alphaeus Dmonte, Roland Oruche, Marcos Zampieri, Prasad Calyam, and Isabelle Augenstein. 2024. Claim verification in the age of large language models: A survey. *arXiv preprint arXiv:2408.14317*.
- Subhabrata Dutta, Rudra Dhar, Prantik Guha, Arpan Murmu, and Dipankar Das. 2023. [A multilingual dataset for identification of factual claims in indian twitter](#). In *Proceedings of the 14th Annual Meeting of the Forum for Information Retrieval Evaluation, FIRE '22*, page 88–92, New York, NY, USA. Association for Computing Machinery.
- Revanth Gangi Reddy, Sai Chetan Chinthakindi, Zhenhailong Wang, Yi Fung, Kathryn Conger, Ahmed Elsayed, Martha Palmer, Preslav Nakov, Eduard Hovy, Kevin Small, and Heng Ji. 2022. [NewsClaims: A new benchmark for claim detection from news with attribute knowledge](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6002–6018, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the association for computational linguistics*, 10:178–206.
- Aaron Halfaker and R. Stuart Geiger. 2020. [Ores: Lowering barriers with participatory machine learning in wikipedia](#). *Preprint*, arXiv:1909.05189.
- Aida Halitaj and Arkaitz Zubiaga. 2024. Providing citations to support fact-checking: Contextualizing detection of sentences needing citation on small wikipedias. *Natural Language Processing Journal*, 8:100093.
- Aida Halitaj and Arkaitz Zubiaga. 2025. Alpet: Active few-shot learning for citation worthiness detection in low-resource wikipedia languages. *Expert Systems With Applications*, 281:127503.
- Maram Hasanain, Reem Suwaileh, Sanne Weering, Chengkai Li, Tommaso Caselli, Wajdi Zaghoulani, Alberto Barrón-Cedeño, Preslav Nakov, and Firoj Alam. 2024. Overview of the clef-2024 checkthat! lab task 1 on check-worthiness estimation of multigenre content. In *Working Notes of the Conference and Labs of the Evaluation Forum (CLEF 2024)*, CEUR Workshop Proceedings, pages 276–286. CEUR Workshop Proceedings. Publisher Copyright: © 2024 Copyright for this paper by its authors.; 25th Working Notes of the Conference and Labs of the Evaluation Forum, CLEF 2024 ; Conference date: 09-09-2024 Through 12-09-2024.
- Naemul Hassan, Bill Adair, James T Hamilton, Chengkai Li, Mark Tremayne, Jun Yang, and Cong Yu. 2015. The quest to automate fact-checking. In *Proceedings of the 2015 computation+ journalism symposium*. The Computation+ Journalism Symposium New York.
- Isaac Johnson, Lucie-Aimée Kaffee, and Miriam Redi. 2024. Wikimedia data for ai: a review of wikimedia datasets for nlp tasks and ai-assisted editing. *arXiv preprint arXiv:2410.08918*.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the nlp world. *arXiv preprint arXiv:2004.09095*.
- Ashkan Kazemi, Kiran Garimella, Devin Gaffney, and Scott A. Hale. 2021. [Claim matching beyond English to scale global fact-checking](#). 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 4504–4517, Online. Association for Computational Linguistics.
- Lev Konstantinovskiy, Oliver Price, Mevan Babakar, and Arkaitz Zubiaga. 2021. Toward automated factchecking: Developing an annotation schema and benchmark for consistent automated claim detection. *Digital threats: research and practice*, 2(2):1–16.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. [From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online. Association for Computational Linguistics.
- Włodzimierz Lewoniewski, Krzysztof Węcel, and Witold Abramowicz. 2017. Analysis of references across wikipedia languages. In *International Conference on Information and Software Technologies*, pages 561–573. Springer.

- Lujun Li, Lama Sleem, Niccolo' Gentile, Geoffrey Nichil, and Radu State. 2025. [Small language models in the real world: Insights from industrial text classification](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)*, pages 971–982, Vienna, Austria. Association for Computational Linguistics.
- Miaoran Li, Baolin Peng, Michel Galley, Jianfeng Gao, and Zhu Zhang. 2024a. Self-checker: Plug-and-play modules for fact-checking with large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 163–181.
- Yufeng Li, Rrubaa Panchendrarajan, and Arkaitz Zubiaga. 2024b. Factfinders at checkthat! 2024: refining check-worthy statement detection with llms through data pruning. *arXiv preprint arXiv:2406.18297*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Laura Majer and Jan Šnajder. 2024. Claim check-worthiness detection: How well do llms grasp annotation guidelines? *arXiv preprint arXiv:2404.12174*.
- Shervin Minaee, Tomas Mikolov, Narjes Nikzad, Meysam Chenaghlu, Richard Socher, Xavier Amatriain, and Jianfeng Gao. 2025. [Large language models: A survey](#). *Preprint*, arXiv:2402.06196.
- Preslav Nakov, David Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barrón-Cedeño, Paolo Papotti, Shaden Shaar, and Giovanni Da San Martino. 2021. Automated fact-checking for assisting human fact-checkers. *arXiv preprint arXiv:2103.07769*.
- OpenAI. 2025. Gpt-5 system card. <https://cdn.openai.com/gpt-5-system-card.pdf>. Accessed: 2025-12-12.
- Rrubaa Panchendrarajan and Arkaitz Zubiaga. 2024. [Claim detection for automated fact-checking: A survey on monolingual, multilingual and cross-lingual research](#). *Natural Language Processing Journal*, 7:100066.
- Subhadarshi Panda and Sarah Ita Levitan. 2021. Detecting multilingual covid-19 misinformation on social media via contextualized embeddings. In *Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda*, pages 125–129.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Miriam Redi, Besnik Fetahu, Jonathan Morgan, and Dario Taraborelli. 2019. Citation needed: A taxonomy and algorithmic assessment of wikipedia's verifiability. In *The world wide web conference*, pages 1567–1578.
- Shaden Shaar, Firoj Alam, Giovanni Da San Martino, Alex Nikolov, Wajdi Zaghouani, Preslav Nakov, and Anna Feldman. 2021. [Findings of the NLP4IF-2021 shared tasks on fighting the COVID-19 infodemic and censorship detection](#). In *Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda*, pages 82–92, Online. Association for Computational Linguistics.
- Paul Suganthan, Fedor Moiseev, Le Yan, Junru Wu, Jianmo Ni, Jay Han, Imed Zitouni, Enrique Alfonseca, Xuanhui Wang, and Zhe Dong. 2025. Adapting decoder-based language models for diverse encoder downstream tasks. *arXiv preprint arXiv:2503.02656*.
- Ahmet Üstün, Viraat Aryabumi, Zheng Yong, Wei-Yin Ko, Daniel D'souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, and 1 others. 2024. Aya model: An instruction finetuned open-access multilingual language model. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15894–15939.
- Andreas Vlachos and Sebastian Riedel. 2014. Fact checking: Task definition and dataset construction. In *Proceedings of the ACL 2014 workshop on language technologies and computational social science*, pages 18–22.
- Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, Shengyi Huang, Kashif Rasul, and Quentin Galouédec. 2020. Trl: Transformer reinforcement learning. <https://github.com/huggingface/trl>.
- Yuxia Wang, Revanth Gangi Reddy, Zain Muhammad Mujahid, Arnav Arora, Aleksandr Rubashevskii, Jiahui Geng, Osama Mohammed Afzal, Liangming Pan, Nadav Borenstein, Aditya Pillai, Isabelle Augenstein, Iryna Gurevych, and Preslav Nakov. 2024. [Factcheck-bench: Fine-grained evaluation benchmark for automatic fact-checkers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14199–14230, Miami, Florida, USA. Association for Computational Linguistics.
- Greta Warren, Irina Shklovski, and Isabelle Augenstein. 2025. Show me the work: Fact-checkers' requirements for explainable automated fact-checking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–21.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.

Jillian Xu, Dylan Zhou, Vinay Shukla, Yang Yang, Junrui Ruan, Shuhuai Lin, Wenfei Zou, Yinxiao Liu, and Karthik Lakshmanan. 2025. [Dual-head reasoning distillation: Improving classifier accuracy with train-time-only reasoning](#). *Preprint*, arXiv:2509.21487.

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, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). *Preprint*, arXiv:1904.09675.

Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2018. [Detection and resolution of rumours in social media: A survey](#). *ACM Comput. Surv.*, 51(2).

## A Data

### A.1 Data collection

We collect claims from Featured Articles, which represent the highest-quality articles on Wikipedia.<sup>10</sup> For languages other than German, English, and Russian, where Featured Articles are abundant, we additionally include Good Articles<sup>11</sup> to increase data coverage. Like Featured Articles, Good Articles are required to contain factually accurate and verifiable information, although they are typically less comprehensive.

For all selected articles across languages, we query the MediaWiki API<sup>12</sup> to obtain their full HTML representations. We then parse each article into paragraphs (excluding the lead section) and split paragraphs into individual claims (i.e., sentences) using Stanza.<sup>13</sup> Cleaning primarily involves correcting formatting issues (e.g., spaces before punctuation), removing extraneous whitespace, and discarding claims shorter than 15 characters, among other steps. These procedures are applied consistently to all retrieved elements, including the topic, the claim itself, and the preceding and following sentences.

Because artifacts introduced during parsing or preprocessing may create spurious correlations and bias our results, we additionally conduct a manual data quality check.

### A.2 Data Labelling

We largely follow prior work in labeling claims (Redi et al., 2019; Halitaj and Zubiaga, 2024), with one minor modification. Consistent with Redi et al. (2019), we label a claim as not requiring a citation only if it appears in a paragraph that contains no citations. Conversely, if a target claim does not contain an inline citation but appears in a paragraph where other claims are cited, we label it as citation needed. This design choice is motivated by our observation that, in practice, claims without inline citations are often implicitly supported by citations in adjacent sentences, either preceding or following the claim, even when no explicit inline citation is present. We provide a qualitative analysis of the general limitations of this heuristic on the German Wikipedia in Appendix C.

<sup>10</sup>[https://en.wikipedia.org/wiki/Wikipedia:Featured\\_articles](https://en.wikipedia.org/wiki/Wikipedia:Featured_articles)

<sup>11</sup>[https://en.wikipedia.org/wiki/Wikipedia:Good\\_articles](https://en.wikipedia.org/wiki/Wikipedia:Good_articles)

<sup>12</sup><https://www.mediawiki.org/wiki/MediaWiki>

<sup>13</sup><https://stanfordnlp.github.io/stanza/>

### A.3 Descriptive Statistics

**Dataset examples** Table 7 illustrates two representative examples in English and Italian from our dataset. In the English examples, the first claim contains multiple factual assertions (e.g., the naming of a city and the reason for its name) related to a historical event, each of which requires a citation to verify that these events occurred. The second English example represents a canonical case of a claim that does not require a citation under Wikipedia’s verifiability policy, as readers can reasonably infer that the source of the information is the film itself.<sup>14</sup>

For the Italian examples, the first claim translates to “In the fall, he began recording a new album, and thus, eleven months after *Due anni dopo*, *L’isola non trovata* was released.” This claim refers to the recording and release of a new album and therefore requires a reliable external source to substantiate the information. The second Italian example introduces a citation to Gianluca Magi. While the cited text itself requires a citation, the surrounding sentence does not, illustrating a case where the claim is not citation-worthy.

**Descriptive statistics** Figure 4 presents descriptive statistics on syntactic distance (a), cross-lingual topic similarity (b), and the ten most frequent topics across languages (c). Our language selection covers a broad range of linguistic diversity, as measured by syntactic distance derived from lang2vec vectors.<sup>15</sup> We cover syntactic distance relationships ranging from distant language pairs (e.g., Indonesian and Albanian) to closely related ones (e.g., Ukrainian and Russian).

Panel (b) illustrates cross-lingual topic similarity measured via cosine similarity between topic vectors. Overall, Wikipedia language editions exhibit relatively high topical overlap. A notable exception is Albanian (*sq*), for which topic similarity is lower, likely due to the substantially smaller number of high-quality Albanian Wikipedia articles in our sample.

Finally, Panel (c) shows the ten most frequent topics across languages, indicating that claims span a diverse set of Wikipedia categories, including Regions, STEM, and Media.

<sup>14</sup>[https://en.wikipedia.org/wiki/Wikipedia:When\\_to\\_cite](https://en.wikipedia.org/wiki/Wikipedia:When_to_cite)

<sup>15</sup><https://github.com/antonisa/lang2vec>

Language	Title	Section	Claim	Previous Sentence	Subsequent Sentence	Label
en	Spanish conquest of Guatemala	San Marcos Province of Tecusitlán and Lacandón	De León renamed the city as San Pedro Sacatepéquez in honour of his friar, Pedro de Angulo.	De León marched to a Maya city named Quezalli by his Nahuatl-speaking allies with a force of fifty Spaniards; his Mexican allies also referred to the city by the name Sacatepequez.	The Spanish founded a village nearby at Candacuchex in April that year, renaming it as San Marcos.	1
	Aliens (film)	Plot	Ripley meets the Colonial Marines aboard the spaceship Sulaco but distrusts their android, Bishop, because the Nostromo's android, Ash, betrayed its crew to protect the alien on company orders.	Still traumatized by her alien encounter, she agrees on the condition that they exterminate the creatures.	-	0
it	Francesco Guccini	Il debutto (1967–1971)	In autunno iniziò le registrazioni di un nuovo disco, e così a undici mesi da Due anni dopo fu pubblicato L'isola non trovata.	-	Il titolo dell'album, che è anche quello di una canzone, è un riferimento a Guido Gozzano; altra citazione letteraria presente nel disco fu quella di J. D. Salinger in La collina.	1
	Alchimia	Alchimia indiana	Gianluca Magi nota come:	L'alchimia giocò un ruolo di spicco fin dalle origini del pensiero indiano.	<citation>	0

Table 7: Dataset Examples.




Figure 4: Descriptive statistics.

#### A.4 Manual Quality Evaluation of MCN

Language	Valid Claim (%)	Label Accuracy (%)	Valid Context (%)
<b>High Resource</b>			
en	0.93	1.00	0.93
pt	1.00	1.00	0.97
de	0.97	1.00	0.93
ru	0.93	0.90	0.93
it	1.00	1.00	1.00
vi	1.00	1.00	0.93
tr	0.93	0.93	0.93
nl	1.00	1.00	1.00
<b>Medium Resource</b>			
uk	1.00	0.97	1.00
ro	0.93	0.97	0.93
id	0.97	1.00	0.97
bg	0.90	1.00	0.97
uz	1.00	1.00	1.00
<b>Low Resource</b>			
no	1.00	1.00	1.00
az	0.97	0.97	0.90
mk	1.00	1.00	1.00
hy	1.00	1.00	1.00
sq	0.93	1.00	0.83

Table 8: Manual data quality evaluation. We assess 540 data (30 per language) examples for the correctness of the claim, its label, and the surrounding context.

To ensure the quality of our data, we conduct a manual evaluation to assess the correctness of the

labels, extracted claims, and context on a representative sample of our data. We randomly sample 540 instances, balanced across languages and labels, and manually evaluate the following: (1) whether the claim is correctly extracted and free from parsing errors (e.g., truncated sentences); (2) whether the assigned citation label is accurate based on our heuristic (see above); and (3) whether the surrounding context is correctly retrieved and free from parsing errors. We conduct these checks by identifying the original claim and its surrounding context on the source Wikipedia page and then manually verifying whether the original instances match our retrieved instances. If they do not match, we mark the corresponding metric as incorrect.

Table 8 presents the data quality evaluation results. Overall, data quality is consistently high across all three criteria and across language resource levels. Most common sources of errors are missing conclusion of certain HTML tags which lead to incorrect claims and labels, and incorrect sentence splitting. However, we also observe discrepancies between the labeling heuristic, also used

in prior work (Redi et al., 2019; Halitaj and Zubiaga, 2024), and real-world editorial citation practices. We discuss this issue in more detail below.

### A.5 Translation Quality Evaluation

Language	BLEU	ROUGE-1	ROUGE-2	BERT Score
<i>Medium-resource</i>				
uk	50.97	0.62	0.36	0.97
ro	63.05	0.62	0.29	0.98
id	58.31	0.57	0.33	0.98
bg	58.64	0.62	0.29	0.97
uz	39.87	0.59	0.27	0.96
<i>Low-resource</i>				
no	63.92	0.71	0.67	0.98
az	42.31	0.53	0.31	0.96
mk	48.93	0.53	0.29	0.94
hy	46.31	0.67	0.31	0.97
sq	57.90	0.37	0.14	0.97

Table 9: Backtranslation quality evaluation.

We conduct an extensive automatic evaluation of the translated parallel test sets used for zero-shot X-CND. Specifically, we back-translate claims that were translated from English into the target languages back into English in order to compute standard translation quality metrics. Automatic evaluation via back-translation is a common procedure when gold reference translations are unavailable (Chatzikoumi, 2020). We use the DeepL API for back-translation.<sup>16</sup> We report standard machine translation metrics, including BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) as n-gram overlap measures, and BERTScore (Zhang et al., 2020) to assess semantic similarity.

Table 9 presents the results of the automatic translation quality evaluation. Across all languages, all metrics indicate strong semantic preservation, with BERTScore values ranging from 0.94 to 0.98. BLEU and ROUGE scores likewise suggest successful reconstruction, although they exhibit greater variance. In particular, Macedonian and Uzbek show lower BLEU and ROUGE scores, likely reflecting greater typological distance from English rather than translation errors, as semantic similarity remains high. Overall, these results provide strong evidence that translation quality is high across languages of all resource levels.

## B Experimental Details

We design all experiments to be runnable on a single A100 GPU (40 GB). All experiments are implemented using PyTorch and the Hugging Face

<sup>16</sup><https://www.deepl.com/en/pro-api>

transformers library<sup>17</sup>. While each experiment fits on a single GPU, we run experiments on an HPC cluster with multiple GPUs to accelerate training.

### B.1 SLMs

Hyperparameter	Value(s)
Epochs	{1, 2, 3, 4}
Learning rate	{ $5 \times 10^{-4}$ , $2 \times 10^{-4}$ , $5 \times 10^{-5}$ }
Batch size	16
Max gradient norm	{0.5, 1.0}
Weight decay	0.01
Seed	42
LoRA rank ( $r$ )	16
LoRA scaling ( $\alpha$ )	32
LoRA dropout	0.05

Table 10: Hyperparameter grid configuration for SFT.

We use SFT (Wei et al., 2021) with QLoRA (Detmeters et al., 2023) to adapt SLMs to CND. For each language–model combination, we perform a hyperparameter grid search using the hyperparameters shown in Table 10. In cases where we encounter out-of-memory errors, we reduce the batch size to 8 and use a gradient accumulation step of 2. We use the `peft`<sup>18</sup> library to implement QLoRA.

Preliminary experiments with different prompts revealed only marginal performance differences. Therefore, we restrict our experiments to a single prompt. We use the simple instruction prompt shown below for all experiments involving SLMs:

```
Instruct Prompt
-----
You are a multilingual classifier.
Decide whether the given {lang} claim
requires a citation.
Use 1 if the claim needs a citation. Use
  0 if the claim does not need a
  citation.

Return only JSON in the format: {"label
": 0} or {"label": 1}.
No explanations or extra text.
```

To enhance the reliability of our results, we select the optimal hyperparameter configuration and repeat training twice using different random seeds. We report the mean performance and standard deviations across these runs.

<sup>17</sup><https://github.com/huggingface/transformers>

<sup>18</sup><https://github.com/huggingface/peft>

## B.2 PLMs

Hyperparameter	Value(s)
Epochs	1,2,3
Learning rate	$\{5 \times 10^{-5}, 1 \times 10^{-5}, 5 \times 10^{-6}\}$
Batch size	$\{16, 32\}$
Max gradient norm	1.0
Weight decay	0.01
Seed	42

Table 11: Hyperparameter grid configuration for PLMs.

For PLMs, we fully fine-tune each model-language combination with hyperparameter grid search using the hyperparameters displayed in Table 11. We run reliability checks identical to the SLM setup.

## B.3 LLMs

For LLMs, we use the same INSTRUCT prompt applied to SLMs as well as a more verbose prompt, as [Majer and Šnajder \(2024\)](#) show that prompt verbosity can influence performance. The verbose prompt provides additional task instructions and incorporates the citation reason taxonomy proposed by [Redi et al. \(2019\)](#):

```

Verbose Prompt
-----
You are a multilingual Wikipedia
citation classifier.
You are provided with a {lang} claim and
its context.
Your task is to analyze the claim and
the context to decide whether the
claim needs a citation.
On Wikipedia, there are various reasons
why a claim may or may not require a
citation.
The reasons are listed below:

Reasons why citations are needed (Label
1):
Quotation      The statement is a direct
quotation or close paraphrase of a
source.
Statistics      The statement contains
statistics or quantitative data.
Controversial   The statement makes
surprising or potentially
controversial claims.
Opinion         The statement expresses a
persons subjective opinion or
belief.
Private Life    The statement contains
claims about a persons private
life.
Scientific      The statement includes
technical or scientific claims.
Historical      The statement makes
general or historical claims that
are not common knowledge.
Other (Needs Citation) The statement
requires a citation for other
reasons.

```

```

Reasons why citations are not needed (
Label 0):
Common Knowledge  The statement
contains well-known or widely
established facts.
Plot              The statement describes the
plot or characters of a work that is
the subject of the article.
Other (No Citation Needed) The
statement does not require a
citation for other reasons.

```

```

Based on these reasons, think step-by-
step to decide in which category the
claim falls.
Return only JSON in the format: {"label
": 0} or {"label": 1}.
No extra text.

```

We use OpenRouter<sup>19</sup> to run inference with the selected LLMs. All LLM experiments are conducted with a low temperature (0.1) which is why we do not perform additional reliability runs.

For few-shot prompting, we randomly sample two examples per label. As shown in Table 4, the inclusion of random shots often improves performance, particularly for the INSTRUCT prompt variant.

## C Results

In this section, we report the results of our main experiments using F1 score instead of accuracy. We further present additional analyses on the effect of different learning rates for few-shot tuning of SLMs.

### C.1 Experiment 1: SFT Variants

Table 12 presents the results of our SFT variant experiments using F1 score as the main evaluation metric.

### C.2 Experiment 2: Multilingual CND Benchmark

Table 13 presents the multilingual CND benchmarking results for all LM families, using F1 score as the main evaluation metric.

### Error Analysis of the German Wikipedia

We observe comparatively low evaluation scores for the German Wikipedia. To better understand this effect, we manually analyze model misclassifications. We find that real-world citation practices of Wikipedia editors occasionally diverge from the labeling heuristic we adopt, which is also used in prior work ([Redi et al., 2019](#); [Halitaj and Zubiaga,](#)

<sup>19</sup><https://openrouter.ai/>

Model	SFT	High Resource										Medium Resource						Low Resource					
		en	pt	de	ru	it	vi	tr	nl	Avg	uk	ro	id	bg	uz	Avg	no	az	mk	hy	sq	Avg	
Aya-Sb	FTL	85.39 (±1.05)	69.63 (±3.94)	61.29 (±1.64)	70.89 (±3.56)	67.20 (±1.56)	73.53 (±0.31)	76.08 (±0.31)	<b>66.76</b> (±2.10)	71.35 (±7.24)	66.72 (±4.84)	70.84 (±2.10)	73.13 (±1.52)	<b>67.20</b> (±0.20)	66.79 (±0.21)	68.94 (±2.91)	64.02 (±3.14)	67.72 (±8.31)	61.53 (±3.44)	54.45 (±1.60)	69.02 (±5.79)	<b>63.35</b> (±3.34)	
	TOL	84.03 (±1.59)	74.48 (±0.99)	<b>66.39</b> (±1.64)	76.03 (±0.99)	70.78 (±0.81)	79.17 (±0.99)	<b>78.50</b> (±0.99)	66.76 (±0.99)	74.52 (±0.99)	60.05 (±0.99)	78.27 (±0.99)	68.41 (±0.99)	66.90 (±0.99)	66.67 (±0.99)	68.06 (±0.99)	<b>67.03</b> (±0.99)	68.69 (±0.99)	<b>66.55</b> (±0.99)	62.35 (±0.99)	66.12 (±0.99)	<b>66.15</b> (±0.99)	
	ES	<b>89.42</b> (±0.42)	<b>81.45</b> (±0.34)	65.57 (±2.89)	<b>77.45</b> (±0.42)	<b>75.45</b> (±0.72)	<b>82.17</b> (±0.72)	<b>62.88</b> (±2.11)	<b>76.40</b> (±1.69)	<b>70.39</b> (±0.71)	<b>79.49</b> (±0.59)	<b>75.41</b> (±1.53)	<b>64.86</b> (±1.19)	<b>69.30</b> (±0.67)	<b>71.89</b> (±0.60)	<b>64.67</b> (±1.17)	<b>72.12</b> (±2.41)	<b>66.58</b> (±1.03)	<b>65.07</b> (±1.12)	<b>73.80</b> (±1.12)	<b>67.65</b> (±1.12)	<b>68.78</b> (±1.12)	
Llama3-8B	FTL	88.93 (±1.05)	79.79 (±1.05)	62.58 (±4.81)	<b>78.85</b> (±0.71)	75.20 (±0.25)	76.38 (±1.61)	<b>62.60</b> (±0.21)	75.68 (±0.96)	69.68 (±1.27)	76.75 (±2.72)	<b>74.08</b> (±0.83)	66.52 (±1.65)	<b>72.65</b> (±2.67)	62.40 (±4.41)	69.29 (±4.23)	63.24 (±2.78)	50.31 (±1.10)	71.83 (±2.78)	71.83 (±2.78)	<b>63.41</b> (±3.84)		
	TOL	89.42 (±1.59)	79.45 (±0.99)	64.85 (±1.64)	77.97 (±0.99)	<b>75.39</b> (±0.99)	<b>77.44</b> (±0.99)	61.76 (±0.99)	<b>76.10</b> (±0.99)	70.28 (±0.99)	75.03 (±0.99)	77.47 (±0.99)	65.76 (±0.99)	69.67 (±0.99)	71.64 (±0.99)	64.34 (±0.99)	69.06 (±0.99)	64.38 (±0.99)	49.80 (±0.99)	71.00 (±0.99)	<b>63.72</b> (±0.99)		
	ES	<b>89.87</b> (±0.20)	<b>80.75</b> (±0.88)	<b>66.00</b> (±2.97)	78.34 (±1.56)	73.32 (±2.01)	<b>82.23</b> (±0.86)	76.61 (±1.03)	58.06 (±1.41)	75.65 (±1.03)	<b>72.38</b> (±1.48)	<b>69.50</b> (±0.99)	<b>77.85</b> (±1.48)	<b>66.76</b> (±1.01)	69.61 (±2.21)	71.22 (±0.63)	<b>66.37</b> (±2.85)	<b>72.12</b> (±1.01)	<b>66.58</b> (±1.01)	<b>65.07</b> (±1.01)	<b>73.76</b> (±1.01)	<b>68.78</b> (±1.01)	
Qwen3-8B	FTL	88.26 (±2.07)	79.92 (±0.97)	58.39 (±6.59)	75.58 (±0.86)	72.98 (±2.18)	81.69 (±0.93)	77.09 (±0.48)	60.86 (±0.76)	74.35 (±10.18)	62.25 (±2.67)	76.71 (±1.30)	75.01 (±0.85)	<b>65.77</b> (±0.89)	70.12 (±0.43)	69.97 (±6.08)	61.69 (±1.56)	71.15 (±2.10)	59.22 (±4.02)	59.58 (±1.19)	70.87 (±1.32)	64.50 (±6.02)	
	TOL	89.24 (±1.49)	<b>80.10</b> (±5.27)	62.97 (±0.83)	<b>78.02</b> (±1.75)	73.16 (±1.99)	79.72 (±1.99)	<b>77.96</b> (±0.99)	61.37 (±2.59)	75.32 (±0.27)	68.54 (±1.75)	75.71 (±2.09)	63.18 (±0.97)	<b>77.20</b> (±1.50)	71.78 (±5.82)	63.61 (±2.18)	68.12 (±2.94)	59.95 (±1.45)	62.95 (±1.38)	<b>73.66</b> (±1.38)	<b>65.66</b> (±1.38)		
	ES	<b>89.25</b> (±0.69)	<b>79.55</b> (±0.05)	<b>66.75</b> (±1.01)	77.75 (±0.70)	<b>75.02</b> (±2.00)	<b>82.82</b> (±0.88)	<b>77.17</b> (±0.91)	<b>63.12</b> (±2.93)	<b>76.43</b> (±8.37)	<b>70.87</b> (±2.69)	<b>78.03</b> (±1.56)	<b>75.80</b> (±2.69)	64.71 (±1.83)	70.03 (±2.10)	<b>71.89</b> (±2.58)	<b>68.26</b> (±1.63)	<b>72.74</b> (±1.34)	<b>63.29</b> (±1.95)	<b>67.02</b> (±1.95)	<b>70.48</b> (±1.98)	<b>68.36</b> (±3.28)	

Table 12: Monolingual CND F1 scores of SLMs across SFT variants (FTL=full-token loss, TOL=target-only loss, ES=encoder-style). **Bold** denotes the best score per model–language pair. Parentheses indicate standard deviations.

		High Resource										Medium Resource						Low Resource					
		en	pt	de	ru	it	vi	tr	nl	Avg	uk	ro	id	bg	uz	Avg	no	az	mk	hy	sq	Avg	Avg
Decoder-based LLMs																							
Zero-shot																							
GPT-5 (Instruct)	79.61	70.06	<u>68.07</u>	73.56	70.99	73.11	75.63	66.30	<u>72.17</u> (±4.20)	67.41	73.76	67.51	67.50	68.78	<u>68.99</u> (±2.79)	66.67	69.54	<b>68.93</b>	66.03	65.49	<u>67.33</u> (±1.80)	69.94 (±3.82)	
GPT-5 (Verbose)	83.36	69.33	66.57	72.11	68.54	74.39	76.06	<u>66.30</u>	<u>72.08</u> (±5.70)	67.13	73.54	68.92	67.68	68.59	<u>69.17</u> (±2.50)	67.49	71.13	67.50	66.76	66.30	<u>67.84</u> (±1.91)	70.09 (±4.43)	
Gemini 2.5 (Instruct)	77.96	69.73	65.51	72.55	71.08	72.29	70.62	63.14	<u>70.36</u> (±4.52)	67.43	72.23	66.16	68.89	69.71	<u>68.69</u> (±2.26)	65.15	66.37	67.06	65.65	65.20	<u>65.89</u> (±0.59)	68.65 (±3.70)	
Gemini 2.5 (Verbose)	81.09	70.85	66.17	71.04	66.67	74.17	74.43	64.81	<u>71.15</u> (±5.40)	67.62	71.52	68.48	68.75	68.18	<u>68.91</u> (±1.52)	66.28	68.29	66.95	65.29	64.75	<u>66.31</u> (±1.40)	69.19 (±4.15)	
Few-shot																							
GPT-5 (Instruct)	84.04	70.14	67.42	74.44	74.43	75.82	77.02	<b>66.67</b>	<u>73.75</u> (±5.07)	66.28	74.51	69.14	<u>69.35</u>	<u>73.68</u>	<u>70.59</u> (±3.43)	66.67	70.52	<u>67.73</u>	66.67	65.54	<u>67.42</u> (±1.90)	71.12 (±4.52)	
GPT-5 (Verbose)	83.04	69.84	65.83	71.49	69.15	75.62	75.74	<b>66.67</b>	<u>72.18</u> (±5.71)	67.41	73.81	69.70	67.60	69.10	<u>69.53</u> (±2.59)	67.04	70.96	67.23	66.94	65.45	<u>67.52</u> (±2.05)	70.15 (±4.48)	
Gemini 2.5 (Instruct)	80.87	69.74	<b>68.64</b>	69.55	69.58	73.54	66.38	62.81	<u>70.14</u> (±5.31)	64.42	47.35	68.57	<b>71.14</b>	69.96	<u>64.29</u> (±3.80)	61.72	63.44	65.27	58.59	66.54	<u>63.11</u> (±3.12)	66.56 (±6.90)	
Gemini 2.5 (Verbose)	81.90	71.67	64.47	71.34	66.67	75.30	72.35	62.88	<u>70.82</u> (±6.18)	64.08	67.13	70.89	68.18	70.26	<u>68.11</u> (±2.72)	65.78	68.11	66.67	66.08	63.92	<u>66.11</u> (±1.52)	68.76 (±4.71)	
Encoder-based PLMs																							
XLM-R-base	88.38	<u>77.47</u>	62.67	<u>77.04</u>	<u>70.56</u>	76.47	<u>77.03</u>	64.51	<u>74.27</u> (±8.22)	67.38	72.46	69.52	64.45	66.67	<u>68.10</u> (±3.43)	66.41	66.61	64.90	<b>67.76</b>	64.64	<u>66.06</u> (±1.97)	70.32 (±6.87)	
XLM-R-large	<b>90.06</b>	<u>79.34</u>	<u>52.23</u>	<u>77.42</u>	<u>74.59</u>	<u>79.86</u>	<u>75.80</u>	44.61	<u>71.74</u> (±9.27)	67.09	<b>79.76</b>	71.18	60.45	66.67	<u>69.03</u> (±2.59)	<b>70.28</b>	69.16	65.79	64.08	66.54	<u>67.17</u> (±3.12)	69.69 (±3.98)	
mBert	89.26	<u>80.87</u>	62.13	75.91	73.52	79.07	74.70	63.04	<u>74.81</u> (±8.99)	69.72	79.11	<b>75.89</b>	63.92	<b>75.49</b>	<b>72.83</b> (±3.01)	67.83	<u>72.13</u> (±0.58)	64.89	66.30	<b>73.85</b>	<b>69.00</b> (±2.31)	72.66 (±3.84)	
Decoder-based SLMs																							
Aya-Sb (ES)	89.42	<b>81.45</b>	65.57	<u>77.45</u>	<b>75.45</b>	82.17	76.83	62.88	<u>76.40</u> (±8.71)	70.39	79.49	75.41	64.86	69.30	71.89 (±0.59)	64.67 (±1.17)	71.63 (±2.41)	65.03 (±1.03)	63.12 (±1.12)	73.80 (±1.12)	67.65 (±1.12)	72.72 (±1.12)	
Llama3-8B (ES)	<b>89.87</b>	<b>80.75</b>	66.00	<b>78.34</b>	73.32	<b>82.23</b>	76.61	58.06	<u>75.65</u> (±1.03)	<b>72.38</b>	69.50	<b>77.85</b>	66.76	69.61	71.22 (±0.63)	<b>66.37</b>	72.12 (±1.03)	66.58 (±1.03)	65.07 (±1.03)	73.76 (±1.03)	<b>68.78</b> (±1.03)		
Qwen3-8B (ES)	89.25	79.55	66.75	<u>77.25</u>	<u>75.02</u>	<b>82.82</b>	<b>77.17</b>	63.12	<b>76.43</b> (±8.37)	70.87 (±2.69)	78.03 (±1.56)	75.80 (±2.69)	64.71 (±1.83)	70.03 (±2.10)	71.89 (±2.58)	<b>68.26</b> (±1.63)	<b>72.74</b> (±1.34)	<b>63.29</b> (±1.95)	<b>67.02</b> (±1.95)	<b>70.48</b> (±1.98)	<b>68.36</b> (±3.28)		

Table 13: Monolingual CND F1 scores across LM families (ES=encoder-style; Verbose and Instruct denote prompt types). **Bold** and underlined scores mark the best and second-best results, respectively. Parentheses indicate standard deviation.

2024). We illustrate these discrepancies with two representative examples.

**Example 1** The first deviation arises from citations that appear only at the end of multiple consecutive paragraphs. In the German Wikipedia article on the city of "Augsburg, the section "Gewässer" consists of six paragraphs. While the first five paragraphs contain no citations, the final paragraph includes a single citation that appears to support the entire multi-paragraph. We observe several such cases, suggesting that editors sometimes choose to include only one citation to cover a broader textual span.

Under the established labeling heuristic, claims in all but the final paragraph would be incorrectly labeled as not requiring a citation. One potential remedy would be to treat all claims preceding the cited paragraph as citation-worthy. However, this assumption does not always hold, as the final citation does not necessarily reference all preceding paragraphs in every multi-paragraph setting.

**Example 2** A second issue arises when a citation attached to an earlier claim implicitly supports subsequent claims without being repeated, provided that the connection is obvious. For example, in the article *Useless*, section *Video and Vinyl*:

2004 veröffentlichte das Bochumer Label Dirty Faces Records das Album *Useless* als Schallplatte und zusätzlich auf farbig gesprenkeltem Vinyl in limitierter Auflage von 525 Stück. [18] Eine weitere Ausgabe des Albums in pinkfarbenem Vinyl erschien im Jahr 2017 bei Drumming Monkey Records.

(English translation) In 2004, the Bochum-based label Dirty Faces Records released the album *Useless* on vinyl and additionally on colored splatter vinyl in a limited edition of 525 copies. [18] Another edition of the album on pink vinyl was released in 2017 by Drumming Monkey Records.

In this context, the citation [18] clearly also supports the final sentence. To account for this editorial practice, we modify existing labeling heuristics (Redi et al., 2019; Halitaj and Zubiaga, 2024) to also label the final sentence as citation needed, which is correct in this case. However, this assumption does not hold for all sentences without explicit citations.

Real-world editorial citation practices are difficult to capture automatically, as they vary across Wikipedia language communities. Overall, these examples highlight the limitations of current heuristic labeling approaches for CND on Wikipedia. Designing more robust labeling strategies therefore remains an important direction for future research.

### C.3 Experiment 3: X-CND

#### Zero-shot X-CND




Figure 5: Zero-shot X-CND F1 scores by test data setting, averaged across language resource groups.

Figure 5 presents zero-shot X-CND results with F1 score as the main metric.

#### Few-shot performance by learning rates

For SLMs fine-tuned with QLoRA, the optimal learning rate is typically relatively higher (2e-4) compared to most other fine-tuning tasks (Dettmers et al., 2023). This is also the case for CND, where we observe that across nearly all model-language combinations, hyperparameter search consistently identifies high learning rates as optimal. However, when applying continued fine-tuning for few-shot transfer following (Lauscher et al., 2020), we quickly observe that this strategy is detrimental to the performance of SLMs.

Therefore, for the second-stage fine-tuning on the target language shots, we treat the learning rate as the sole hyperparameter to optimize. While we could also include other hyperparameters, the connection between high learning rates, small amounts of training shots, and poor model performance seems most obvious to remedy.

Model	LR	Shots				Avg $\Delta$
		k = 50	k = 100	k = 250	k = 500	
<b>Medium Resource</b>						
Llama 3-8B (FTL)	1e-04	66.57 ( $\pm 4.64$ )	66.24 ( $\pm 4.83$ )	65.19 ( $\pm 5.98$ )	66.78 ( $\pm 5.35$ )	-0.88
	1e-05	67.20 ( $\pm 5.19$ )	67.60 ( $\pm 5.39$ )	67.53 ( $\pm 5.73$ )	67.78 ( $\pm 5.39$ )	<b>0.46</b>
	1e-06	66.91 ( $\pm 5.33$ )	67.58 ( $\pm 5.27$ )	67.38 ( $\pm 5.18$ )	68.02 ( $\pm 5.59$ )	0.40
	1e-07	66.91 ( $\pm 5.42$ )	67.47 ( $\pm 5.21$ )	67.23 ( $\pm 5.27$ )	68.07 ( $\pm 5.41$ )	0.35
Llama 3-8B (TOL)	1e-04	60.90 ( $\pm 4.81$ )	60.52 ( $\pm 7.95$ )	66.08 ( $\pm 9.82$ )	70.03 ( $\pm 8.31$ )	-2.06
	1e-05	65.83 ( $\pm 6.13$ )	66.81 ( $\pm 5.99$ )	65.14 ( $\pm 6.63$ )	65.95 ( $\pm 5.47$ )	-0.52
	1e-06	66.32 ( $\pm 5.92$ )	66.82 ( $\pm 5.99$ )	66.28 ( $\pm 5.30$ )	67.26 ( $\pm 5.42$ )	0.23
	1e-07	66.59 ( $\pm 5.84$ )	67.03 ( $\pm 5.89$ )	66.64 ( $\pm 5.49$ )	67.10 ( $\pm 6.35$ )	<b>0.39</b>
Llama 3-8B (ES)	1e-04	65.55 ( $\pm 6.58$ )	64.07 ( $\pm 5.90$ )	69.13 ( $\pm 6.00$ )	72.38 ( $\pm 5.73$ )	-0.41
	1e-05	69.36 ( $\pm 4.45$ )	69.80 ( $\pm 4.44$ )	70.24 ( $\pm 4.63$ )	71.14 ( $\pm 4.88$ )	<b>1.95</b>
	1e-06	68.39 ( $\pm 4.56$ )	68.67 ( $\pm 4.54$ )	69.23 ( $\pm 4.46$ )	69.84 ( $\pm 4.26$ )	0.84
	1e-07	68.19 ( $\pm 4.88$ )	68.25 ( $\pm 4.84$ )	68.33 ( $\pm 4.88$ )	68.43 ( $\pm 4.63$ )	0.11
mBert	1e-04	55.39 ( $\pm 7.49$ )	59.83 ( $\pm 9.79$ )	59.62 ( $\pm 7.34$ )	58.11 ( $\pm 7.21$ )	-2.25
	5e-05	62.18 ( $\pm 7.42$ )	63.33 ( $\pm 7.32$ )	64.77 ( $\pm 7.50$ )	63.74 ( $\pm 6.89$ )	<b>3.37</b>
	1e-05	60.96 ( $\pm 5.93$ )	61.31 ( $\pm 6.02$ )	61.52 ( $\pm 5.40$ )	62.63 ( $\pm 5.75$ )	1.13
	1e-06	60.45 ( $\pm 6.17$ )	60.52 ( $\pm 6.17$ )	60.48 ( $\pm 6.14$ )	60.44 ( $\pm 6.11$ )	-0.01
1e-07	60.49 ( $\pm 6.11$ )	60.51 ( $\pm 6.14$ )	60.47 ( $\pm 6.21$ )	60.49 ( $\pm 6.14$ )	0.01	
<b>Low Resource</b>						
Llama 3-8B (FTL)	1e-04	62.61 ( $\pm 4.17$ )	61.88 ( $\pm 3.96$ )	61.42 ( $\pm 5.12$ )	62.71 ( $\pm 4.97$ )	-0.44
	1e-05	62.92 ( $\pm 4.12$ )	62.35 ( $\pm 3.30$ )	62.60 ( $\pm 4.35$ )	63.72 ( $\pm 4.39$ )	0.31
	1e-06	62.90 ( $\pm 4.12$ )	62.47 ( $\pm 3.30$ )	62.44 ( $\pm 4.35$ )	63.65 ( $\pm 4.39$ )	0.27
	1e-07	62.82 ( $\pm 3.93$ )	62.50 ( $\pm 3.43$ )	62.38 ( $\pm 4.28$ )	63.97 ( $\pm 4.58$ )	<b>0.32</b>
Llama 3-8B (TOL)	1e-04	57.76 ( $\pm 4.73$ )	58.43 ( $\pm 5.02$ )	60.74 ( $\pm 6.08$ )	63.22 ( $\pm 5.57$ )	-0.61
	1e-05	60.96 ( $\pm 5.24$ )	60.88 ( $\pm 5.13$ )	60.10 ( $\pm 4.72$ )	61.31 ( $\pm 4.50$ )	0.06
	1e-06	60.58 ( $\pm 4.97$ )	60.88 ( $\pm 4.45$ )	61.16 ( $\pm 3.91$ )	61.69 ( $\pm 3.89$ )	<b>0.42</b>
	1e-07	61.19 ( $\pm 3.90$ )	60.66 ( $\pm 6.07$ )	60.53 ( $\pm 4.82$ )	61.49 ( $\pm 4.69$ )	0.31
Llama 3-8B (ES)	1e-04	59.41 ( $\pm 4.29$ )	58.25 ( $\pm 5.23$ )	62.44 ( $\pm 6.33$ )	64.40 ( $\pm 5.21$ )	-1.17
	1e-05	63.31 ( $\pm 4.21$ )	62.98 ( $\pm 4.66$ )	63.47 ( $\pm 5.05$ )	64.43 ( $\pm 5.20$ )	<b>1.25</b>
	1e-06	62.68 ( $\pm 5.59$ )	62.79 ( $\pm 5.61$ )	62.96 ( $\pm 5.56$ )	63.85 ( $\pm 4.93$ )	0.77
	1e-07	62.30 ( $\pm 5.64$ )	62.34 ( $\pm 5.62$ )	62.56 ( $\pm 5.52$ )	62.52 ( $\pm 5.51$ )	0.13
mBert	1e-04	53.92 ( $\pm 6.44$ )	55.84 ( $\pm 7.25$ )	57.31 ( $\pm 7.45$ )	58.08 ( $\pm 8.44$ )	-0.62
	5e-05	58.97 ( $\pm 6.77$ )	60.27 ( $\pm 6.87$ )	59.76 ( $\pm 4.93$ )	62.16 ( $\pm 6.50$ )	<b>3.37</b>
	1e-05	57.09 ( $\pm 4.42$ )	57.37 ( $\pm 4.33$ )	57.49 ( $\pm 4.10$ )	58.04 ( $\pm 4.23$ )	0.58
	1e-06	56.93 ( $\pm 4.47$ )	56.90 ( $\pm 4.46$ )	56.93 ( $\pm 4.42$ )	56.93 ( $\pm 4.42$ )	0.01
1e-07	56.87 ( $\pm 4.45$ )	56.90 ( $\pm 4.47$ )	56.83 ( $\pm 4.43$ )	56.91 ( $\pm 4.47$ )	-0.03	

Table 14: Few-shot accuracy for X-CND. LR=learning rate. Parentheses indicate standard deviations. We observe that generative SFT yields consistent improvements only when using substantially reduced learning rates, whereas mBERT and ES-based tuning remain effective under higher learning rates.

Tables 14 and 15 report the results of our experiments analyzing the effect of learning rate choice on few-shot X-CND. We observe a consistent pattern: Llama (ES) and mBERT benefit from relatively higher learning rates, whereas generative SFT variants (FTL and TOL) exhibit improvements from few-shot tuning only when using substantially lower learning rates, down to 1e-7. Even then, the gains achieved by generative SFT through additional few-shot fine-tuning remain modest. Overall, these results indicate that learning rate selection is critical for continued fine-tuning on small amounts of target-language data for QLoRA-based SLMs, particularly under generative SFT. In contrast, mBERT performs optimally with learning

Model	LR	Shots				Avg $\Delta$				
		k = 50	k = 100	k = 250	k = 500					
<b>Medium Resource</b>										
Llama 3-8B (FTL)	1e-04	67.41 ( $\pm 3.41$ )	(-2.85)	64.82 ( $\pm 4.34$ )	(-5.45)	65.07 ( $\pm 6.59$ )	(-5.20)	66.67 ( $\pm 5.47$ )	(-3.60)	-4.27
	1e-05	70.19 ( $\pm 3.98$ )	(-0.07)	70.35 ( $\pm 4.14$ )	(+0.09)	69.61 ( $\pm 4.66$ )	(-0.66)	69.06 ( $\pm 4.57$ )	(-1.20)	-0.46
	1e-06	70.16 ( $\pm 4.07$ )	(-0.11)	70.70 ( $\pm 3.67$ )	(+0.43)	70.34 ( $\pm 4.14$ )	(+0.08)	70.93 ( $\pm 4.22$ )	(+0.66)	0.27
	1e-07	70.16 ( $\pm 4.15$ )	(-0.10)	70.65 ( $\pm 3.81$ )	(+0.38)	70.31 ( $\pm 4.17$ )	(+0.05)	71.18 ( $\pm 3.95$ )	(+0.92)	<b>0.31</b>
Llama 3-8B (TOL)	1e-04	63.68 ( $\pm 6.51$ )	(-6.49)	60.20 ( $\pm 7.88$ )	(-9.97)	67.72 ( $\pm 8.67$ )	(-2.46)	71.60 ( $\pm 7.58$ )	(+1.42)	-4.37
	1e-05	67.71 ( $\pm 4.59$ )	(-2.47)	67.81 ( $\pm 4.36$ )	(-2.36)	65.49 ( $\pm 5.13$ )	(-4.69)	66.40 ( $\pm 5.47$ )	(-3.78)	-3.32
	1e-06	69.95 ( $\pm 4.20$ )	(-0.22)	70.31 ( $\pm 4.43$ )	(+0.13)	69.45 ( $\pm 3.92$ )	(-0.72)	69.36 ( $\pm 4.18$ )	(-0.82)	-0.41
	1e-07	69.77 ( $\pm 4.65$ )	(-0.40)	70.13 ( $\pm 4.82$ )	(+0.05)	69.97 ( $\pm 4.21$ )	(-0.20)	70.28 ( $\pm 4.61$ )	(+0.10)	<b>-0.14</b>
Llama 3-8B (ES)	1e-04	69.06 ( $\pm 4.44$ )	(+1.72)	63.22 ( $\pm 11.34$ )	(-4.12)	69.40 ( $\pm 6.72$ )	(+2.06)	73.46 ( $\pm 5.80$ )	(+6.11)	1.44
	1e-05	70.67 ( $\pm 3.82$ )	(+3.32)	70.49 ( $\pm 4.32$ )	(+3.15)	71.03 ( $\pm 4.99$ )	(+3.69)	71.90 ( $\pm 4.48$ )	(+4.55)	<b>3.68</b>
	1e-06	67.84 ( $\pm 7.85$ )	(+0.50)	68.56 ( $\pm 6.69$ )	(+1.22)	69.40 ( $\pm 5.53$ )	(+2.06)	70.11 ( $\pm 4.40$ )	(+2.77)	1.64
	1e-07	67.36 ( $\pm 8.87$ )	(+0.02)	67.48 ( $\pm 8.74$ )	(+0.14)	67.60 ( $\pm 8.56$ )	(+0.26)	67.76 ( $\pm 8.14$ )	(+0.42)	0.21
mBert	1e-04	51.78 ( $\pm 28.71$ )	(-2.19)	59.87 ( $\pm 21.01$ )	(+5.90)	56.43 ( $\pm 21.12$ )	(+2.47)	48.66 ( $\pm 28.22$ )	(-5.30)	0.22
	5e-05	63.05 ( $\pm 16.79$ )	(+9.09)	66.79 ( $\pm 9.08$ )	(+12.82)	69.34 ( $\pm 5.88$ )	(+15.38)	65.41 ( $\pm 12.03$ )	(+11.45)	<b>12.19</b>
	1e-05	55.84 ( $\pm 15.93$ )	(+1.87)	56.60 ( $\pm 15.21$ )	(+2.63)	57.17 ( $\pm 13.88$ )	(+3.21)	59.28 ( $\pm 12.47$ )	(+5.32)	3.26
	1e-06	53.94 ( $\pm 17.70$ )	(-0.03)	53.99 ( $\pm 17.72$ )	(+0.03)	53.99 ( $\pm 17.65$ )	(+0.03)	53.89 ( $\pm 17.77$ )	(-0.08)	-0.01
1e-07	53.93 ( $\pm 17.76$ )	(-0.03)	53.95 ( $\pm 17.77$ )	(-0.01)	53.94 ( $\pm 17.78$ )	(-0.03)	53.99 ( $\pm 17.74$ )	(+0.02)	-0.01	
<b>Low Resource</b>										
Llama 3-8B (FTL)	1e-04	61.21 ( $\pm 5.43$ )	(-2.74)	58.89 ( $\pm 6.74$ )	(-5.06)	57.50 ( $\pm 6.23$ )	(-6.44)	58.82 ( $\pm 6.57$ )	(-5.12)	-4.84
	1e-05	64.22 ( $\pm 4.49$ )	(+0.28)	63.15 ( $\pm 4.50$ )	(-0.79)	62.45 ( $\pm 4.48$ )	(-1.49)	62.57 ( $\pm 4.73$ )	(-1.38)	-0.85
	1e-06	64.50 ( $\pm 4.80$ )	(+0.56)	63.88 ( $\pm 4.35$ )	(+0.06)	63.40 ( $\pm 5.46$ )	(-0.55)	65.09 ( $\pm 4.45$ )	(+1.15)	0.28
	1e-07	64.49 ( $\pm 4.83$ )	(+0.54)	63.96 ( $\pm 4.43$ )	(+0.02)	63.47 ( $\pm 5.34$ )	(-0.47)	65.64 ( $\pm 5.12$ )	(+1.69)	<b>0.45</b>
Llama 3-8B (TOL)	1e-04	56.74 ( $\pm 7.51$ )	(-5.78)	55.48 ( $\pm 7.30$ )	(-7.05)	54.39 ( $\pm 13.34$ )	(-8.14)	59.38 ( $\pm 9.49$ )	(-3.14)	-6.02
	1e-05	59.54 ( $\pm 7.81$ )	(-2.98)	56.44 ( $\pm 7.51$ )	(-6.08)	54.72 ( $\pm 6.46$ )	(-7.80)	56.80 ( $\pm 7.39$ )	(-5.72)	-5.65
	1e-06	62.89 ( $\pm 7.09$ )	(+0.37)	62.58 ( $\pm 6.44$ )	(+0.06)	61.70 ( $\pm 6.15$ )	(-0.82)	61.25 ( $\pm 5.62$ )	(-1.27)	-0.41
	1e-07	62.40 ( $\pm 8.03$ )	(-0.12)	61.82 ( $\pm 8.67$ )	(-0.70)	61.52 ( $\pm 7.05$ )	(-1.00)	62.85 ( $\pm 5.33$ )	(+0.33)	<b>-0.37</b>
Llama 3-8B (ES)	1e-04	61.01 ( $\pm 8.95$ )	(-0.26)	51.65 ( $\pm 16.46$ )	(-9.62)	61.91 ( $\pm 6.45$ )	(+0.64)	63.98 ( $\pm 6.84$ )	(+2.71)	-1.63
	1e-05	58.39 ( $\pm 6.97$ )	(-2.88)	61.48 ( $\pm 5.37$ )	(+0.21)	62.42 ( $\pm 4.94$ )	(+1.15)	64.29 ( $\pm 4.16$ )	(+3.02)	0.37
	1e-06	61.28 ( $\pm 10.65$ )	(+0.01)	61.80 ( $\pm 9.36$ )	(+0.53)	61.57 ( $\pm 7.68$ )	(+0.30)	62.42 ( $\pm 4.67$ )	(+1.15)	<b>0.50</b>
	1e-07	61.14 ( $\pm 11.91$ )	(-0.13)	61.31 ( $\pm 11.56$ )	(+0.04)	61.54 ( $\pm 11.12$ )	(+0.27)	61.50 ( $\pm 10.80$ )	(+0.22)	0.10
mBert	1e-04	46.34 ( $\pm 30.89$ )	(-3.79)	44.50 ( $\pm 29.78$ )	(-5.64)	55.81 ( $\pm 18.73$ )	(+5.67)	55.14 ( $\pm 22.42$ )	(+5.00)	0.31
	5e-05	57.78 ( $\pm 17.88$ )	(+7.64)	60.39 ( $\pm 10.91$ )	(+10.25)	62.30 ( $\pm 7.50$ )	(+12.16)	65.13 ( $\pm 6.03$ )	(+14.99)	<b>11.26</b>
	1e-05	51.42 ( $\pm 17.46$ )	(+1.28)	52.17 ( $\pm 15.91$ )	(+2.03)	53.17 ( $\pm 14.55$ )	(+3.03)	55.03 ( $\pm 12.63$ )	(+4.89)	2.81
	1e-06	50.15 ( $\pm 18.85$ )	(+0.01)	50.15 ( $\pm 18.80$ )	(+0.01)	50.20 ( $\pm 18.76$ )	(+0.06)	50.21 ( $\pm 18.74$ )	(+0.07)	0.04
1e-07	50.10 ( $\pm 18.83$ )	(-0.04)	50.10 ( $\pm 18.89$ )	(-0.04)	49.99 ( $\pm 18.91$ )	(-0.15)	50.14 ( $\pm 18.80$ )	(0.00)	-0.06	

Table 15: Few-shot F1 scores for X-CND. LR=learning rate. Parentheses indicate standard deviations.

rates similar to those identified during hyperparameter.

### Few-shot Performance

Table 16 reports the results for few-shot CND with F1 score as the main metric on non-parallel data.

### Generalizability

Non-Parallel Data													
Model	Shots	Medium-Resource					$\Delta$ Avg	Low-Resource					$\Delta$ Avg
		uk	ro	id	bg	uz		no	az	mk	hy	sq	
Llama3-8B (FTL)	0	65.47 ( $\pm 1.25$ )	75.55 ( $\pm 1.87$ )	71.91 ( $\pm 0.53$ )	65.87 ( $\pm 2.76$ )	72.52 ( $\pm 1.92$ )	-0.10	65.03 ( $\pm 0.39$ )	67.03 ( $\pm 2.54$ )	63.11 ( $\pm 0.39$ )	56.39 ( $\pm 3.19$ )	68.16 ( $\pm 1.73$ )	+0.54
	$\Delta$ 50	+0.27 ( $\pm 2.18$ )	+0.21 ( $\pm 1.76$ )	-0.47 ( $\pm 1.23$ )	+0.25 ( $\pm 0.74$ )	-0.76 ( $\pm 1.80$ )	-0.44 ( $\pm 1.33$ )	+0.14 ( $\pm 1.21$ )	+2.04 ( $\pm 2.26$ )	+0.71 ( $\pm 5.59$ )	+0.27 ( $\pm 2.70$ )	+0.27 ( $\pm 2.70$ )	+0.02
	$\Delta$ 100	-0.15 ( $\pm 1.04$ )	-0.84 ( $\pm 0.43$ )	+0.56 ( $\pm 0.54$ )	+1.61 ( $\pm 1.17$ )	+0.74 ( $\pm 1.12$ )	+0.38 ( $\pm 1.54$ )	-0.93 ( $\pm 1.23$ )	+0.88 ( $\pm 2.50$ )	+0.21 ( $\pm 3.83$ )	-0.76 ( $\pm 1.34$ )	-0.76 ( $\pm 1.34$ )	+0.02
	$\Delta$ 250	-0.62 ( $\pm 1.24$ )	+0.05 ( $\pm 1.40$ )	+1.08 ( $\pm 0.77$ )	+1.35 ( $\pm 1.07$ )	-1.64 ( $\pm 2.01$ )	+0.05 ( $\pm 1.82$ )	-0.64 ( $\pm 1.76$ )	-1.50 ( $\pm 2.06$ )	-0.69 ( $\pm 6.26$ )	-0.09 ( $\pm 1.73$ )	-0.09 ( $\pm 1.73$ )	-0.47
	$\Delta$ 500	+0.80 ( $\pm 0.45$ )	+0.23 ( $\pm 1.50$ )	+0.64 ( $\pm 1.03$ )	+1.55 ( $\pm 0.56$ )	+1.37 ( $\pm 1.65$ )	+0.92 ( $\pm 1.97$ )	+0.14 ( $\pm 2.51$ )	+0.82 ( $\pm 2.42$ )	+6.80 ( $\pm 4.06$ )	+0.64 ( $\pm 1.77$ )	+0.64 ( $\pm 1.77$ )	+1.69
Llama3-8B (TOL)	0	65.33 ( $\pm 1.53$ )	75.20 ( $\pm 1.76$ )	72.70 ( $\pm 0.83$ )	66.21 ( $\pm 2.53$ )	71.43 ( $\pm 2.66$ )	-0.40	64.91 ( $\pm 1.33$ )	67.44 ( $\pm 3.26$ )	63.57 ( $\pm 2.85$ )	50.53 ( $\pm 1.88$ )	66.16 ( $\pm 1.68$ )	-0.12
	$\Delta$ 50	-0.25 ( $\pm 2.70$ )	+0.07 ( $\pm 1.94$ )	-1.38 ( $\pm 2.80$ )	-0.36 ( $\pm 3.99$ )	-0.09 ( $\pm 3.88$ )	-1.33 ( $\pm 0.52$ )	-1.33 ( $\pm 3.80$ )	+1.82 ( $\pm 3.21$ )	+0.05 ( $\pm 13.86$ )	+0.20 ( $\pm 2.22$ )	+0.20 ( $\pm 2.22$ )	-0.12
	$\Delta$ 100	-0.26 ( $\pm 0.89$ )	+0.10 ( $\pm 3.64$ )	-0.34 ( $\pm 1.74$ )	-0.21 ( $\pm 5.46$ )	+0.46 ( $\pm 2.59$ )	-0.05 ( $\pm 3.28$ )	+0.44 ( $\pm 2.46$ )	-0.68 ( $\pm 2.85$ )	-4.26 ( $\pm 1.74$ )	+1.56 ( $\pm 2.56$ )	+1.56 ( $\pm 2.56$ )	-0.70
	$\Delta$ 250	+0.84 ( $\pm 1.60$ )	+1.03 ( $\pm 1.26$ )	-1.07 ( $\pm 0.66$ )	+0.08 ( $\pm 2.77$ )	-1.90 ( $\pm 2.68$ )	-0.20 ( $\pm 4.83$ )	-2.52 ( $\pm 4.45$ )	-0.63 ( $\pm 2.15$ )	-0.54 ( $\pm 4.61$ )	+1.69 ( $\pm 2.04$ )	+1.69 ( $\pm 2.04$ )	-1.00
	$\Delta$ 500	+0.07 ( $\pm 1.33$ )	+0.56 ( $\pm 1.72$ )	-0.02 ( $\pm 0.35$ )	-0.85 ( $\pm 3.39$ )	+0.76 ( $\pm 0.17$ )	+0.10 ( $\pm 3.64$ )	-1.24 ( $\pm 3.05$ )	-0.93 ( $\pm 4.11$ )	+5.60 ( $\pm 7.38$ )	+1.32 ( $\pm 2.04$ )	+1.32 ( $\pm 2.04$ )	+0.33
Llama3-8B (ES)	0	61.00 ( $\pm 6.35$ )	74.93 ( $\pm 1.79$ )	74.08 ( $\pm 1.81$ )	56.86 ( $\pm 10.82$ )	69.84 ( $\pm 2.15$ )	+3.32	66.28 ( $\pm 1.45$ )	68.44 ( $\pm 1.34$ )	65.27 ( $\pm 2.46$ )	39.54 ( $\pm 8.23$ )	66.83 ( $\pm 2.44$ )	+0.01
	$\Delta$ 50	+6.24 ( $\pm 0.43$ )	+0.68 ( $\pm 2.37$ )	-2.00 ( $\pm 2.56$ )	+10.35 ( $\pm 2.17$ )	+1.34 ( $\pm 3.06$ )	+3.15	-0.98 ( $\pm 1.88$ )	-0.63 ( $\pm 0.64$ )	-0.41 ( $\pm 2.51$ )	+2.05 ( $\pm 6.48$ )	+0.05 ( $\pm 2.37$ )	+0.01
	$\Delta$ 100	+6.71 ( $\pm 0.27$ )	-0.59 ( $\pm 1.73$ )	-0.13 ( $\pm 2.30$ )	+8.41 ( $\pm 4.63$ )	+1.35 ( $\pm 2.72$ )	+3.15	-1.06 ( $\pm 0.96$ )	+0.18 ( $\pm 0.97$ )	-1.07 ( $\pm 1.36$ )	+5.03 ( $\pm 5.33$ )	-0.45 ( $\pm 2.85$ )	+0.53
	$\Delta$ 250	+4.96 ( $\pm 0.91$ )	+0.84 ( $\pm 0.82$ )	+1.87 ( $\pm 1.77$ )	+8.68 ( $\pm 2.29$ )	+2.11 ( $\pm 3.01$ )	+3.69	-2.66 ( $\pm 0.76$ )	-0.74 ( $\pm 0.91$ )	-2.09 ( $\pm 0.58$ )	+8.17 ( $\pm 5.01$ )	-1.18 ( $\pm 2.68$ )	+0.30
	$\Delta$ 500	+7.36 ( $\pm 0.71$ )	-0.58 ( $\pm 2.05$ )	+1.64 ( $\pm 1.46$ )	+8.96 ( $\pm 1.91$ )	+5.39 ( $\pm 2.87$ )	+4.55	-3.91 ( $\pm 1.31$ )	-0.69 ( $\pm 1.73$ )	-3.52 ( $\pm 0.78$ )	+15.59 ( $\pm 2.65$ )	-1.70 ( $\pm 2.53$ )	+1.15
mBert	0	54.97 ( $\pm 14.70$ )	59.43 ( $\pm 16.59$ )	59.98 ( $\pm 16.82$ )	40.36 ( $\pm 25.63$ )	55.07 ( $\pm 19.87$ )	+9.09	54.95 ( $\pm 15.31$ )	45.93 ( $\pm 23.11$ )	48.12 ( $\pm 22.58$ )	46.23 ( $\pm 23.77$ )	55.47 ( $\pm 22.38$ )	+7.64
	$\Delta$ 50	+2.51 ( $\pm 11.86$ )	+15.14 ( $\pm 2.55$ )	+10.39 ( $\pm 1.52$ )	+9.35 ( $\pm 31.00$ )	+8.07 ( $\pm 16.39$ )	+9.09	-5.00 ( $\pm 26.22$ )	+20.95 ( $\pm 3.40$ )	+1.90 ( $\pm 25.16$ )	+7.76 ( $\pm 17.85$ )	+12.59 ( $\pm 0.72$ )	+7.64
	$\Delta$ 100	+4.83 ( $\pm 11.24$ )	+13.43 ( $\pm 10.41$ )	+10.55 ( $\pm 1.72$ )	+20.87 ( $\pm 11.09$ )	+14.43 ( $\pm 3.07$ )	+12.82	-1.09 ( $\pm 15.98$ )	+19.95 ( $\pm 3.96$ )	+11.37 ( $\pm 11.18$ )	+8.74 ( $\pm 13.87$ )	+12.30 ( $\pm 2.79$ )	+10.25
	$\Delta$ 250	+7.92 ( $\pm 5.88$ )	+16.75 ( $\pm 2.74$ )	+12.26 ( $\pm 1.24$ )	+26.58 ( $\pm 3.78$ )	+13.39 ( $\pm 5.37$ )	+15.38	-0.28 ( $\pm 8.65$ )	+22.43 ( $\pm 0.36$ )	+14.44 ( $\pm 8.40$ )	+12.45 ( $\pm 7.04$ )	+11.74 ( $\pm 0.79$ )	+12.16
	$\Delta$ 500	-1.26 ( $\pm 19.98$ )	+9.20 ( $\pm 11.08$ )	+10.79 ( $\pm 0.75$ )	+21.57 ( $\pm 8.71$ )	+16.95 ( $\pm 8.25$ )	+11.45	+9.46 ( $\pm 3.41$ )	+22.45 ( $\pm 2.73$ )	+10.18 ( $\pm 11.00$ )	+19.77 ( $\pm 2.77$ )	+13.09 ( $\pm 1.10$ )	+14.99

Table 16: Few-shot CND F1 scores on non-parallel data by model and number of shots.  $\Delta$  denotes gains over zero-shot accuracy. Parentheses report standard deviations. All models are fine-tuned on English claims only.

	en	it	uk	ro	no	az
Random articles	70.40	65.49	71.43	67.27	67.19	55.77
Random articles (corrected labels)	91.93	80.00	83.23	79.49	93.26	70.20
FA baseline	89.87	73.32	72.38	69.50	66.37	72.12
RA Distribution (0/1)	14/85	11/89	6/92	7/92	8/92	2/97

Table 17: Monolingual CND F1 scores on the Random Articles (RA) test set ( $n = 100$  per language) using Llama 3 (ES). Due to under-referencing in random articles, we manually corrected the labels.