

Cross-Dialect Information Retrieval: Information Access in Low-Resource and High-Variance Languages

Robert Litschko^{1,2} Oliver Kraus¹ Verena Blaschke^{1,2} Barbara Plank^{1,2}

¹ MaiNLP, Center for Information and Language Processing, LMU Munich, Germany

² Munich Center for Machine Learning (MCML), Munich, Germany

{robert.litschko, o.kraus2, verena.blaschke, b.plank}@lmu.de

Abstract

A large amount of local and culture-specific knowledge (e.g., people, traditions, food) can only be found in documents written in dialects. While there has been extensive research conducted on cross-lingual information retrieval (CLIR), the field of cross-dialect retrieval (CDIR) has received limited attention. Dialect retrieval poses unique challenges due to the limited availability of resources to train retrieval models and the high variability in non-standardized languages. We study these challenges on the example of German dialects and introduce the first German dialect retrieval dataset, dubbed WikiDIR, which consists of seven German dialects extracted from Wikipedia. Using WikiDIR, we demonstrate the weakness of lexical methods in dealing with high lexical variation in dialects. We further show that the commonly used zero-shot cross-lingual transfer approach with multilingual encoders does not transfer well to extremely low-resource setups, motivating the need for resource-lean and dialect-specific retrieval models. We finally demonstrate that (document) translation is an effective way to reduce the dialect gap in CDIR.

1 Introduction

Cross-lingual information retrieval (CLIR) is the task of retrieving documents written in a language different from the query language. The promise of CLIR is to make information accessible across language boundaries. One of its main challenges lies in bridging the *lexical gap between languages* (Berger et al., 2000), which is caused by the fact that different languages use different vocabularies. This problem has been extensively studied in the context of ad-hoc news retrieval (Braschler, 2004; Lawrie et al., 2022) and Wikipedia passage retrieval (Sasaki et al., 2018; Sun and Duh, 2020; Ogundepo et al., 2022; Li et al., 2022), among others.

Relevant information may, however, not only be




Figure 1: Approximate locations of German regional languages and dialects studied in this work: Low German (nds), North Frisian (frr), Saterfrisian (stq), Ripuarian (ksh), Rhine Franconian (pfl), Alemannic (als), and Bavarian (bar). Image based on a map of Europe by Marian Sigler, CC BY-SA 3.0.

written in a different language but also in a regional dialect or language variant. This is particularly true for culturally-related concepts such as local traditions, individuals, and locations. In this work, we focus on cross-dialect information retrieval (CDIR) with German regional languages and dialects (see Figure 1). The *lexical gap between dialects* in CDIR is caused by orthographic variations and different regional expressions, resulting from the lack of standardization. For example, the German term *München* (“Munich”) is also spelled as *Münche* (als), *Minga* (bar), *Münsche* (ksh) and *Minche* (pfl). In addition to variations found in different dialects, there are also variations that arise from subdialects. The dialect variation renders lexical retrieval methods such as BM25 (Robertson et al., 2009) and CLIR approaches based on machine translation (Lignos et al., 2019; Artetxe et al., 2023) less effective, since they only match a single variation. At the

same time, many dialects only have a small digital footprint. All dialects used in this work are categorized as low-resource languages (Joshi et al., 2020). It is unclear how well resource-lean methods such as zero-shot cross-lingual transfer paradigm based on pre-trained language models (MacAvaney et al., 2020; Shi et al., 2020; Litschko et al., 2022, 2023) or large language models (Sun et al., 2023) perform on CDIR. While dialect variation has recently gained traction in the NLP community (Faisal et al., 2024), it is still understudied in the context of retrieval. The only recent work we are aware of is by Chari et al. (2023), who compare retrieval performance between British English and American English spelling variants found in MS MARCO (Bajaj et al., 2016), and Valentini et al. (2024), who compare retrieval models on different Spanish variants. To close this gap, we introduce WikiDIR, a cross-dialect retrieval dataset based on German dialect Wikipedias. WikiDIR includes queries in standard German and documents in seven regional language variants, which we will refer to as dialects throughout the rest of this paper. Different from prior work, we specifically focus on orthographic and lexical variation on languages that are non-standardized and have limited resources. Common approaches to derive queries from Wikipedia include using article titles (Sun and Duh, 2020; Ogundepo et al., 2022) or the first sentence (Sasaki et al., 2018). These approaches, however, fall short of capturing full spectrum of dialect variation. To address this limitation, and as another contribution, we manually annotated entity mentions in for five dialects and built dictionaries that contain different spelling variants checked by native speakers.

In summary, we are the first to study dialect-specific lexical variation in a low-resource and high-variation setting. We release our WikiDIR dataset, dialect dictionaries, as well as our annotations and source code.¹ We address the following research questions: **RQ1:** How well do lexical and neural retrieval models perform on CDIR? **RQ2:** To what extent can we improve their performance by specializing models for dialects models by means of continual pretraining? **RQ3:** How does the lack of standardization in dialects, i.e., the dialect gap, affect retrieval performance? **RQ4:** How well do large language models fare when employed as dialect translation models?

¹<https://github.com/MaiNLP/WikiDIR>

2 Related Work

Wikipedia-based test collections. Wikipedia has been extensively used as a resource for training and evaluating cross-lingual retrieval models in the context of ad-hoc retrieval (Schamoni et al., 2014; Sasaki et al., 2018; Sun and Duh, 2020; Li et al., 2022; Ogundepo et al., 2022) and cross-lingual question answer retrieval (Roy et al., 2020; Asai et al., 2021). Queries and documents are typically extracted from article titles and article texts. Most CLIR datasets derived from Wikipedia obtain relevance labels by propagating synthetic monolingual relevance labels through inter-language links. They can be broadly grouped by how monolingual labels are derived. One approach is to synthesize relevance labels based on (mutual) links between Wikipedia articles (Schamoni et al., 2014; Sasaki et al., 2018; Frej et al., 2020; Li et al., 2022). The second approach, used in more recent benchmarks such as CLIRMatrix (Sun and Duh, 2020) and AfriCLIRMatrix (Ogundepo et al., 2022), performs monolingual lexical retrieval and obtains relevance labels by discretizing BM25 retrieval scores. In this work, we adopt and extend this approach (cf. Section 3). The use of synthetic queries and relevance labels has been criticized in (Zhang et al., 2023b), because they bias CLIR evaluation towards a lexical notion of relevance, rather than semantic relevance. Other Wikipedia-based CLIR datasets such as Mr.TyDi (Zhang et al., 2021) and XOR-QA (Asai et al., 2021) are centered around QA and require a more sophisticated query interpretation. Our focus, however, lies in *dialect-specific lexical variation* that can be attributed to the lack of standardization.

Cross-lingual transfer. Two common ways to bridge the language gap in CLIR include the so-called translate-train and translate-test approach (Saleh and Pecina, 2020; Bonifacio et al., 2021; Artetxe et al., 2023). The drawback of machine-translated test collections is that they not only suffer from translationese (Zhang and Toral, 2019; Zhao et al., 2020) but are also culturally biased towards the language from which they are translated from (Hecht and Gergle, 2010; Gutiérrez et al., 2016). Due to their low-resource nature, dialects studied in this work are not supported by neural machine translation such as Google Translate or NLLB (Costa-jussà et al., 2022). Unlike German dialects, standard German is regarded as a high-resource language (Joshi et al., 2020). We

	Wikipedia articles (total)				Total	Relevance assessments (avg)			
	#Train	#Dev	#Test	#Analysis		#Train	#Dev	#Test	#Analysis
German	64,253	8,032	8,031	–	80,316	4.8	5.0	2.4	–
Low German (nds)	58,901	7,382	7,367	470	74,120	12.8	12.2	6.5	269.9 / 315.9
Alemannic (als)	17,092	2,121	2,146	4,639	25,998	13.4	13.2	6.9	35.0 / 60.0
Bavarian (bar)	17,974	2,234	2,244	718	13,170	8.9	8.8	4.4	77.1 / 91.6
North Frisian (frr)	12,550	1,594	1,565	–	15,709	9.9	10.6	4.7	–
Saterfrisian (stq)	2,710	342	339	–	3,391	9.5	8.9	4.5	–
Ripuarian (ksh)	1,800	228	222	210	2,460	6.6	6.5	3.4	19.9 / 30.3
Rhine Franconian (pfl)	1,971	244	250	157	2,622	3.6	3.2	1.8	23.0 / 36.8
Total/Average	177,251	22,177	22,164	6,194	227,786	9.2	9.1	4.6	85.0 / 106.3

Table 1: WikiDIR statistics. We report the number of Wikipedia articles, which corresponds to the number of queries and documents. We also show for each dialect the average number of relevance annotations. In addition to the standard train-dev-test splits, we also use a held-out analysis split consisting of queries for which we found at least one dialect-specific lexical variation (see Section 3.2). To quantify the lexical dialect gap, we provide relevance assessments without (✗) / with (✓) documents containing dialect variations.

therefore study the commonly-adopted zero-shot cross-lingual transfer paradigm (Hu et al., 2020; Liang et al., 2020) for CDIR. Specifically, we study the zero-shot transfer based on multilingual BERT (Devlin et al., 2019) and the recently proposed reranking approach based on large language models (LLMs) (Sun et al., 2023).

Dialect retrieval. Most of the dialects studied in this work are often overlooked in prior work. For example, Artemova and Plank (2023) study bilingual lexicon induction (i.e., word-level CDIR) on Bavarian and Alemannic. On sentence-level CDIR, WikiMatrix (Schwenk et al., 2021) covers Low German and Bavarian. Vamvas et al. (2024) investigate continual pretraining for Swiss German sentence retrieval. On passage-level retrieval, Chari et al. (2023) separate American from British spelling variants found in MS MARCO and find that neural retrieval models, despite showing a clear preference towards the former, exhibit robust retrieval performance. Most recently, Valentini et al. (2024) present MessIRve, a retrieval benchmark with queries sourced from different Spanish-speaking regions using Google’s autocomplete API and documents corresponding to Spanish Wikipedia articles. Finally, the CLIRMatrix benchmark (Sun and Duh, 2020) uses English queries and covers three out of four dialects studied in this work (als, bar, nds). Different from prior work, we investigate the CDIR performance in truly low-resource settings and contribute bilingual dictionaries containing diverse lexical and orthographic dialect variants, which allow us to quantify the impact of the dialect gap.

3 WikiDIR Benchmark

Our benchmark includes data from seven regional languages and dialects spoken in Germany: North Frisian (wiki code: *frr*), Saterfrisian (*stq*), Low German (*nds*), Ripuarian (*ksh*), Rhine Franconian (*pfl*), Alemannic (*als*), and Bavarian (*bar*). Most of these wikis contain multiple dialect subgroups.

3.1 Dataset Pipeline

We now describe our dataset pipeline used to create WikiDIR, which closely resembles the approach proposed in CLIRMatrix (Sun and Duh, 2020). Figure 2 shows our dataset pipeline and Table 1 summarizes key statistics for each dialect.

Queries and documents. We use Wikipedia article titles as queries and treat as documents the first 200 extracted tokens from the article texts. We specifically use the display title of each dialect article, which is different from the canonical title written in standard. Since our focus is on lexical variation between German and its dialects, we exclude all articles where the queries consist solely of numbers. Many Wikipedia articles contain lexical shortcuts, i.e., references to the official German spelling in the first sentence: *Minga (amtli: München) [ˈmɪŋ(:)ɐ] is d’Haptstod vo Bayern.*² To address this issue and the risk for the task of becoming a mere (monolingual) lexical keyword matching task, we use German article titles to eliminate such lexical shortcuts. The example thus becomes: *Minga [ˈmɪŋ(:)ɐ] is d’Haptstod vo Bayern.*

²“Munich (officially *München*) is the capital of Bavaria.”

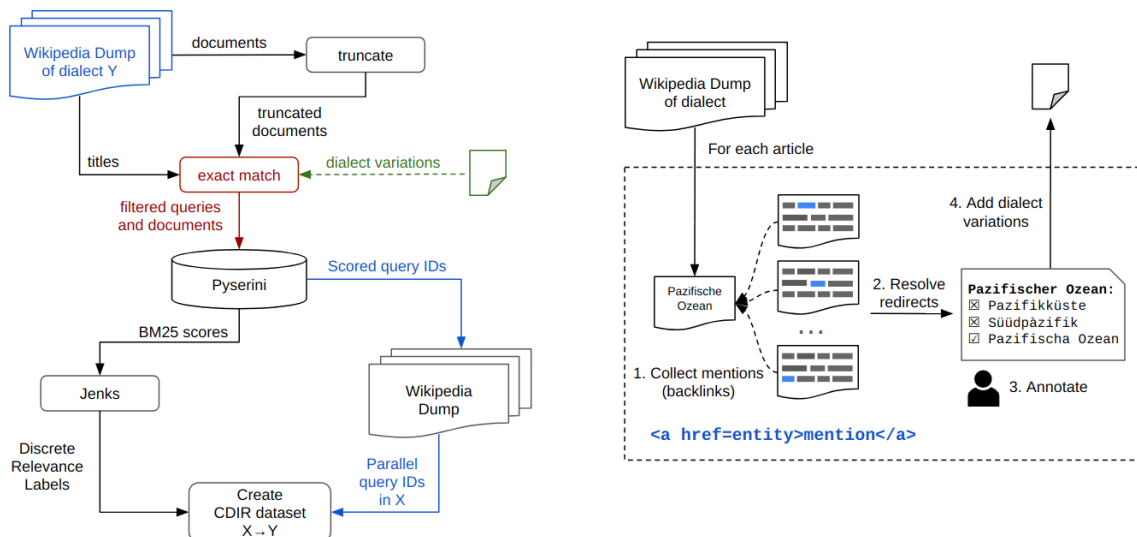


Figure 2: **Left:** Dataset pipeline used to create WikiDIR. Our approach deviates from CLIRMatrix (Sun and Duh, 2020) in three aspects (highlighted in color). First, we only label documents as relevant if they contain an exact match of the query. Second, we use dialect dictionaries to also label documents as relevant if they contain any dialect variation of the query. Third, we use inter-language links to replace dialect queries by their German translation, rather than propagating labels on the document side. **Right:** Annotation pipeline used to build dialect dictionaries containing orthographic and lexical variations. We extract candidates of dialect variations (entity mentions) from incoming links. Correct variations are human-annotated and added to the dictionary. We provide for a set of held-out queries (analysis split) relevance assessments with (✓) and without (✗) documents containing dialect variations.

Relevance labels. Obtaining human-labeled relevance annotations is expensive and does not scale to a large number of languages. Because of this, we follow prior work (Sun and Duh, 2020; Ogundepo et al., 2022) and opt for synthetic relevance labels obtained from lexical retrieval scores. That is, we first obtain monolingual relevance labels from discretized lexical retrieval scores. For this, we use the BM25 implementation provided by the Pyserini framework (Lin et al., 2021). This approach, however, falls short in dealing with multi-term queries containing terms that are identical in the dialect and German. Consider for example the Alemannic query *Kanton Lozärn* (“Canton of Lucerne”). Under the bag-of-words assumption a BM25 retriever would first return many documents containing the German term *Kanton*. Labeling documents with partial lexical matches as relevant would therefore lead to many false positives, which skews the retrieval evaluation towards standard German rather than dialect retrieval. Because of this, we only score documents with BM25 (and thus label as relevant) if they contain an exact phrasal match. To obtain monolingual labels, we follow the approach proposed by Sun and Duh (2020) and use min-max normalization to normalize all BM25 values to the unit range and discretize them into the labels 1–5

using Jenks natural breaks optimization (McMaster and McMaster, 2002). Documents extracted from the same articles as queries are labeled with 6. Finally, to obtain cross-lingual relevance labels, we use inter-language links and replace dialect titles (i.e., dialect queries) with German titles.

3.2 Dialect Variation

Relying on Wikipedia titles to annotate the relevance of documents does not account for the dialect variance resulting from the lack of standardization. In this work, we investigate dialect variation from two perspectives: (1) orthographic and lexical variation and (2) variation within subdialects.

Orthographic and lexical variation. Broadly speaking, dialect variation can be categorized into orthographic differences based on different regional pronunciations (e.g., *Minga* and *Mincha* referring to Munich) and different regional word choices (e.g., *Brötchen*, *Weckli*, and *Semml* referring to bread rolls) (Barbour and Stevenson, 1998). While lexical and morphological variation in German dialects and regional languages have been well-documented, e.g., by Elspaß and Möller (2003), there exist very few publicly available resources. Recent efforts on lexical resources include, e.g., building crowdsourcing translations for Bavar-

Dialect	Anchor links		Dictionary	
	Entities	Mentions	Entities	Variants
nds	1,653	6,579	484	1.55
als	6,440	38,671	4,684	3.11
bar	1,402	6,289	721	2.12
ksh	232	1,056	211	2.92
pfl	182	862	157	2.79
Total/Avg	9,909	53,457	6,257	2.50

Table 2: Number of entities and entity mentions (i.e., candidates) per dialect, as well as the average number of variants per entity after annotation.

ian (Burghardt et al., 2016) and lexicon induction via rule extraction (Millour and Fort, 2019). To study the impact of variation on CDIR, we build entity-centric dialect dictionaries from human annotations on five dialects. The annotation pipeline is illustrated in Figure 2. For each entity, we first extract all backlinks from the raw Wikipedia dumps. The anchor link text typically corresponds to entity mentions written in dialect, while the href attribute stores the standard German spelling. We use the Wikimedia REST API to resolve redirects and obtain normalized dialect article titles. This way, we collect for each dialect title (i.e., query) all mentions. As shown in Table 2, we extracted over 53k mentions for 9.9K entities in five dialects.

Many entity mentions do not correspond to lexical variations in a strict sense. Consider for example the Alemannic dialect title *Schtadt* (“city”), while *Schtatt* and *Sdadd* refer to correct orthographic variations, the term *Schtädt* refers to its plural form (“cities”). Next to grammatical number, we also consider deviations in grammatical gender, grammatical case and word classes. Lexical normalization is typically done through stemming or lemmatization. However, such lexical analyzers do not yet exist for German dialects. Also, entity mentions often correspond to abbreviations or semantically related terms. This is for example the case for the term *Zitig* (“newspaper”), where *Zitung* is a correct variation while the compound noun *Daageszitig* refers to a daily newspaper. Since we do not want to inflate the dialect gap with different preprocessing steps, we only focused on dialect variations in a strict sense (regional word choices and spelling variations), rather than variations that also occur in standard German (different word inflections, synonyms, paraphrases). These variations were annotated by native German speakers

with a computational linguistics background. We calculated the inter-annotator agreement on a sample of three dialects and 50 queries in terms of Cohen’s Kappa (Cohen, 1960). The Kappa values for *pfl* ($\kappa = 0.83$), *bar* ($\kappa = 0.87$), and for *als* ($\kappa = 0.80$) all show high agreements. Notably, the resulting dialect dictionaries vary considerably in size (cf. Table 2). This is to some extent due to the fact that smaller Wikipedias have fewer editors and fewer links between articles (Arora et al., 2024). On average, our dictionaries store for each dialect query 2.5 dialect-specific variations.

We use our dictionaries to create a held-out analysis split consisting of the subset of queries for which we found at least one dialect variation (Figure 2, right). To measure the impact of dialect variations on retrieval (see Tables 4 and 6), we create two relevance assessments for our analysis split. In the first version (✗) we label documents containing the query (i.e., dialect title) as relevant (e.g., Bavarian documents containing the term *Minga*). In the second version (✓), we additionally label all documents that contain any dialect variation as relevant (*Münch’n*, *Minkcha*, *Minkn*, *Minchn*, *Mingna*, *Minkhn*, *Münchn*). From here, we obtain cross-lingual relevance labels as described in Section 3.1. We make all dictionaries, together with WikiDIR, publicly available for future uptake.

Dialect groups. To investigate language variation within individual wikis, we use the subset of articles tagged with fine-grained dialect information. We use the dialect groupings by Wiesinger (1983) and Walker and Wilts (2001) to assign articles (i.e., subdialects) to comparable dialect subgroups. This is possible for Bavarian (North, Central, and South Bavarian), Alemannic (Swabian, and Low, Central, High, and Highest Alemannic), and North Frisian (Insular and Mainland North Frisian). We do not divide Riparian or Rhine Franconian into subgroups, since they each only contain one of Wiesinger’s (1983) groups. There are no dialect tags available for Saterfrisian, and too few for meaningful analysis in the case of Low German.

4 Models

We benchmark retrieval models that are widely-used in monolingual and cross-lingual information retrieval. Using WikiDIR, we investigate how well they perform on CDIR and to what extent they are able to deal with dialect-specific variation.

4.1 Multi-Stage Retrieval

We first investigate the multi-stage retrieval paradigm, which has been proposed by [Nogueira et al. \(2019\)](#). We use BM25 as a first-stage retriever ([Robertson et al., 2009](#)) implemented in the Pyserini framework ([Lin et al., 2021](#)). We follow prior work and re-rank the top 100 documents using MonoBERT ([MacAvaney et al., 2019](#); [Litschko et al., 2022](#)). MonoBERT is a cross-encoder model that treats relevance prediction as a sequence-pair classification task ([Nogueira and Cho, 2019](#)). Specifically, MonoBERT concatenates query document pairs [CLS] Query [SEP] Doc [SEP] and predicts the relevance score of each document independently of other documents. Here, we opt for multilingual BERT (mBERT, [Devlin et al., 2019](#)) as the encoder model. We use the pairwise hinge loss ([MacAvaney et al., 2019](#); [Sun and Duh, 2020](#)) and compare MonoBERT in two settings. We first fine-tune it on the respective training splits of WikiDIR. Given the low-resource nature of many dialects, we are also interested in the zero-shot transfer from German. To this end, we also train a MonoBERT model only on the standard German portion of WikiDIR.

4.2 Re-Ranking with LLMs

[Sun et al. \(2023\)](#) demonstrate the effectiveness of commercial LLMs on zero-shot re-ranking tasks. To this end, an LLM is prompted with the query and the respective retrieved documents. The task of the model is then to generate a new ranking of the documents based on their relatedness to the query. As a large number of retrieved documents can quickly fill up the context window of LLMs, the authors propose a sliding-window approach for re-ranking. This method iteratively re-ranks subsets of documents, moving backwards through the list until all documents have been evaluated. Building on this work, [Zhang et al. \(2023a\)](#) demonstrate the effectiveness of this approach on open-source LLMs. In our work, we therefore adapt this approach to evaluate its effectiveness on low-resource dialect data. For our experiments we use the Meta-Llama-3-8B-Instruct³ model as a state-of-the-art open-source LLM. We use the same prompt template for re-ranking as presented by [Sun et al. \(2023\)](#). We use greedy decoding with a sliding window size of 16 and a step size of 8. This is a scaled-down version

³<https://huggingface.co/meta-llama/Meta-Llama-3.1-8B-Instruct>

```
===== LLM-RERANKING =====

system: You are RankGPT, an intelligent assistant that can
rank passages based on their relevancy to the query.

user: I will provide you with num passages, each indicated
by number identifier []. Rank them based on their relevance
to query: {{query}}.

assistant: Okay, please provide the passages.

user: [1] {{passage_1}}

assistant: Received passage [1]

user: [2] {{passage_2}}

assistant: Received passage [2]

(more passages) ...

user:
Search Query: {{query}}.
Rank the num passages above based on their relevance to
the search query. The passages should be listed in descend-
ing order using identifiers, and the most relevant passages
should be listed first, and the output format should be [] >
[], e.g., [1] > [2]. Only response the ranking results, do not
say any word or explain.
```

of the settings used by [Sun et al. \(2023\)](#), to ensure that we do not exceed the context window of our model. Similar to MonoBERT, we re-rank the top 100 documents retrieved by BM25. If BM25 returns less than 100 documents, they are not filled up with random documents.

4.3 ColBERTv2

Naturally, the performance of multi-stage retrieval models is bound by the recall of the first-stage ranker. [Khattab and Zaharia \(2020\)](#) introduce the ColBERT model, which achieves effective and efficient retrieval with a late interaction mechanism between query and document. Documents and queries are separately encoded, which are then compared using a MaxSim operator, allowing for token-level matching during retrieval. ColBERT can be used for both re-ranking and full retrieval. [Santhanam et al. \(2022\)](#) introduce ColBERTv2, which decreases the stored index size while providing higher quality retrieval results.

We train ColBERTv2 models using mBERT⁴ models. All models are used for full retrieval on

⁴<https://huggingface.co/google-bert/bert-base-multilingual-uncased>

the datasets, rather than re-ranking. Due to our comparatively small datasets, we modify the standard hyperparameters by [Santhanam et al. \(2022\)](#) by reducing the batch size to 16 and not using warm-up steps. Additionally, as with MonoBERT, we train a ColBERT model only on the standard German portion of WikiDIR to investigate its zero-shot cross-lingual transfer performance.

4.4 Continual Pretraining of MLMs

[Senel et al. \(2024\)](#) show that by jointly pretraining a masked language model on a low resource language together with a related high resource language, they can improve cross-lingual transfer. Specifically, they introduce Bilingual Alternating LM-ing (BALM) and Bilingual Joint LM-ing (BJLM). For BALM, the model is trained on monolingual batches from a high-resource and a low-resource language in an alternating fashion. For BJLM, each batch contains an equal ratio of high-resource and low-resource instances. In this work, we pretrain multiple masked language models using BALM and BJLM to investigate whether continual pretraining can improve the retrieval performance for low-resource dialects.

Given that our dataset consists of multiple closely related dialects, we also experiment with Multilingual Alternating LM-ing (MALM) and Multilingual Joint LM-ing (MJLM). For MALM, the model is trained on a batch of German instances, followed by a batch from each of the seven dialects in our dataset, until all instances from the largest dataset have been processed. For MJLM, each batch contains an equal ratio of high-resource samples (i.e., German) and samples from all seven dialects in our dataset. This approach ensures that the model is exposed to a balanced ratio of the high-resource language and all low-resource dialects simultaneously during training.

We continually mBERT using BJLM, BALM, BALM and MALM with the train splits of all seven dialects of WikiDIR. We train all models with a learning rate of $1e-4$ and a batch size of 128 for one epoch. All processed documents are split into chunks of 128 tokens.

4.5 Document Translation with LLMs

[Adeyemi et al. \(2024\)](#) investigate the effectiveness of using LLMs as translation models in the context of CLIR with African languages. In this work, we investigate the suitability of the Meta-Llama-3-8B-Instruct model as a dialect translation model.

We use greedy decoding to translate documents from different dialects to standard German. The translated documents are used (1) as training data for ColBERTv2 and MonoBERT models, and (2) at retrieval time. We follow [Adeyemi et al. \(2024\)](#) and use the following prompt template:

```
===== TRANSLATE TRAIN PROMPT =====
```

```
Document: {document}
```

```
Translate this document to German. Only return the translation, do not say any other word.
```

5 Results

In the following, we first compare different retrieval paradigms on CDIR (RQ1, RQ2), followed by a detailed analysis on the impact of dialect variance (RQ3). We then discuss the effectiveness of the document translation approach (RQ4). We evaluate our results with the normalized discounted cumulative gain of the top-10 ranked items (nDCG@10), as implemented in the `ir_measures` framework ([MacAvaney et al., 2022](#)).

5.1 Main Results

Table 3 shows our main results. We find that BM25 already performs competitively, achieving an average nDCG@10 score of 0.737. Between both neural IR models, we find that only ColBERT, when fine-tuned on the respective training portions of CDIR (ColBERT_{Fine-tuned}), outperforms BM25. When used in the context of zero-shot cross-lingual transfer (MonoBERT_{Zero-shot} and ColBERT_{Zero-shot}), both models perform worse. Similarly, our LLM-based baseline, which also has not been fine-tuned on WikiDIR data, fails to outperform BM25. Overall, we find that on average ColBERT_{Fine-tuned} outperforms other retrieval models (RQ1). We hypothesize that this is due to the fact that its performance is not capped by a (lexical) first-stage retriever. Further specializing retrieval models through bilingual (BJLM, BALM) and multilingual (MJLM, MALM) continual pretraining can bring improvements for the dialects with larger Wikipedias (RQ2).

5.2 Dialect Variation

We now investigate dialect variation on the lexical level. The results reported in Table 4 are computed on the analysis split. As before, queries are written in German and documents are written in different dialects. Compared to our results on the test

Model	<i>als</i>	<i>bar</i>	<i>nds</i>	<i>frr</i>	<i>stq</i>	<i>ksh</i>	<i>pfl</i>	avg
BM25	0.750	0.818	0.905	0.655	0.690	0.669	0.669	0.737
Llama-3 (reranking)	0.741	0.805	0.904	0.647	0.666	0.640	0.662	0.723
MonoBERT _{Zero-shot}	0.581	0.660	0.703	0.552	0.629	0.575	0.623	0.618
MonoBERT _{Fine-tuned}	0.720	0.733	0.781	0.588	0.658	0.649	0.663	0.685
+BJLM	0.691	0.758	0.749	0.617	0.648	0.646	0.646	0.680
+BALM	0.618	0.688	0.792	0.621	0.669	0.675	0.647	0.673
+MJLM	0.593	0.754	0.801	0.606	0.648	0.659	0.659	0.674
+MALM	0.663	0.762	0.669	0.587	0.637	0.635	0.635	0.655
Co1BERT _{Zero-shot}	0.790	0.842	0.890	0.682	0.735	0.628	0.780	0.764
Co1BERT _{Fine-tuned}	0.770	0.859	0.907	0.729	0.753	0.670	0.810	0.785
+BJLM	0.763	0.860	0.911	0.744	0.679	0.545	0.778	0.754
+BALM	0.796	0.878	0.912	0.737	0.707	0.597	0.790	0.774
+MJLM	0.756	0.840	0.906	0.737	0.684	0.560	0.781	0.752
+MALM	0.743	0.851	0.905	0.742	0.661	0.656	0.774	0.762

Table 3: Cross-dialect retrieval results on WikiDIR languages and averages over all languages in terms of nDCG@10. Queries are written in German and documents are written in dialects. **Bold**: Best result for each language.

Model	Variations	<i>als</i>	<i>bar</i>	<i>nds</i>	<i>pfl</i>	<i>ksh</i>	avg
BM25	✗	0.401	0.483	0.371	0.293	0.401	0.390
	✓	0.357	0.453	0.347	0.277	0.325	0.353
MonoBERT	✗	0.407	0.402	0.307	0.285	0.354	0.351
	✓	0.369	0.387	0.291	0.277	0.300	0.325
Co1BERT	✗	0.490	0.469	0.412	0.422	0.349	0.428
	✓	0.475	0.460	0.399	0.385	0.336	0.411

Table 4: Cross-dialect retrieval results on WikiDIR in terms of nDCG@10. We evaluate BM25 and the fine-tuned versions of MonoBERT and Co1BERT on two versions of the WikiDIR analysis split. The first version disregards dialect variations (✗) and considers any document relevant that contains the query. The second version considers any document as relevant if it contains the query or any of its dialect variations (✓).

Subdialect	#Articles	nDCG@10
<i>Alemannic dialect groups</i>		
Swabian	891	0.831
Low Alemannic	3362	0.777
Central Alemannic	221	0.538
High Alemannic	1467	0.800
Highest Alemannic	89	0.547
<i>North Frisian dialect groups</i>		
Insular North Frisian	6980	0.611
Mainland North Frisian	1825	0.485
<i>Bavarian dialect groups</i>		
Central Bavarian	4461	0.599
South Bavarian	297	0.612
North Bavarian	89	0.784

Table 5: CDIR results of BM25 on dialect subgroups.

split, we find that the results are overall substantially lower. This is because titles in dialect articles,⁵ which are known to have different variations, are more likely to be different from their German counterpart. All three models consistently perform worse when we include relevant documents that contain dialect spelling variants (✓). The lexical dialect gap is larger for BM25 and MonoBERT. This is expected, since lexical methods are limited to exact term matches between queries and documents. Furthermore, our neural models have not been explicitly trained to match different dialect-specific spelling variants.

We will now investigate CDIR results where the queries have been grouped into their dialect subgroups, as shown in Table 5. Here, too, we find large differences in CDIR performance between different subgroups. This demonstrates the high dialectal variation among subdialects (**RQ3**).

⁵We used dialect articles to compute monolingual relevance labels (see Section 3.1).

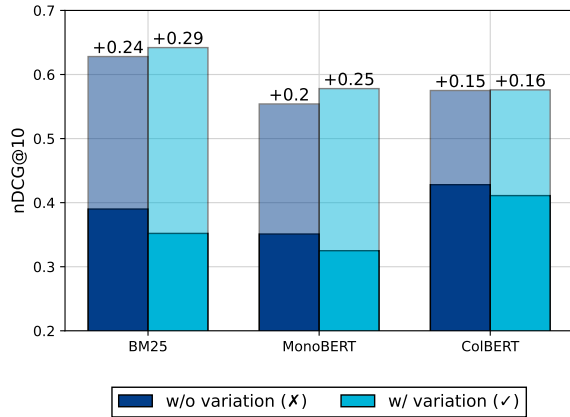


Figure 3: CDIR results on the original training and test data and results (gains) obtained from document translation. We report the average over five dialects.

5.3 Document translation

In this section, we compare the performance between models trained on WikiDIR against models that were trained (and evaluated) on documents, which have been translated into German, as described in Section 4.5. In Figure 3, we compare the average results obtained with documents written in dialects against documents translated to German. We find that document translation yields large gains across all languages (see also Table 6 in Appendix A), with gains ranging from +0.15 to +0.29 nDCG@10. Unsurprisingly, BM25 benefits the most among our three models. This shows that document translation into a language with a standard orthography (German) is an effective way of reducing the dialectal orthographic variance (RQ4).

We now compare the performance with (✓) and without (X) dialect variations. On the original data, we find that dialect variations (✓) cause a drop in performance. However, on the translated documents we observe the opposite. We hypothesize that this related to the larger number of documents with exact matches (see Table 7 in Appendix A). That is, in the absence of the lexical dialect gap, treating more documents as relevant (i.e., increasing the recall base) simplifies the retrieval task.

Our results in Figure 3 and Table 6 suggest that document translation is an effective way to reduce the dialect gap. However, it is important to note that Llama-3 has been exposed to Wikipedia during pretraining. This means that the model has also been exposed to lexical shortcuts (see Section 3.1) and pretrained on dialect code-switched data. Finally, our qualitative analysis reveals two important

shortcomings (see Table 8). First, they indicate that employing Llama-3 for dialect translation tends to be more effective in handling orthographic variations linked to different regional pronunciations and sometimes struggles with lexical variations caused by regional word choices. Also, when combined with BM25, the bag-of-words assumption can lead to many false positives.

6 Conclusion

In this work, we introduce WikiDIR, a cross-dialect information retrieval (CDIR) test collection spanning seven German dialects. We additionally release dialect dictionaries containing different spelling variants, facilitating a detailed analysis of their impact on CDIR. Our results show that there is a substantial gap between retrieving documents that contain the query and documents containing the query or any of its dialect variations. We also show that, even within a specific dialect, there are large differences in retrieval performance among its subdialects. WikiDIR paves the way for future work on retrieval models that generalize across lexical variants as well as dialect-specific lexical analyzers (e.g. stemmers, lemmatizers).

7 Limitations

In this work, we focused on lexical dialect variation. This neglects the semantic aspect and the information asymmetry between dialects and standard German. Furthermore, our approach for extracting dialect variation candidates (Section 3.2), which we then annotate to build dialect dictionaries, relies on the connectivity of Wikipedias (Arora et al., 2024). In future work, we plan to investigate other approaches of extracting dictionaries without relying on links between articles.

Wikipedias in low-resource languages can be more prone to quality issues (Tatariya et al., 2024). Almost all of the wikis included in this study are in the best-quality tier determined by Tatariya et al.’s (2024) set of heuristics (with the exception of *pfl*; tier 3 of 4). However, this might be an overestimation due to the quality heuristics favouring wikis with many unique words (which might be inflated due to spelling variation within dialect wikis; cf. Section 3.2). In terms of the collaborative depth (Meta contributors, 2024), the wikis in this study range mostly in the middle, according to both Wikipedia’s Depth measure and Alshahrani et al.’s (2023) DEPTH+ measure.

Finally, due to resource constraints, we did not fine-tune Llama-3 for dialect reranking.

Acknowledgments

We thank the members of the MaiNLP lab for their insightful feedback on earlier drafts of this paper. We specifically thank Marwin Hättrich for the initial discussion of this work, and Fabian David Schmidt, Philipp Mondorf, Miriam Winkler and Michael A. Hedderich for their help with annotating dialect variations. We thank the anonymous reviewers for their insightful comments. This research is supported by the ERC Consolidator Grant DIALECT 101043235.

References

- Mofetoluwa Adeyemi, Akintunde Oladipo, Ronak Pradeep, and Jimmy Lin. 2024. [Zero-shot cross-lingual reranking with large language models for low-resource languages](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–656, Bangkok, Thailand. Association for Computational Linguistics.
- Saied Alshahrani, Norah Alshahrani, and Jeanna Matthews. 2023. [DEPTH+: An enhanced depth metric for Wikipedia corpora quality](#). In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pages 175–189, Toronto, Canada. Association for Computational Linguistics.
- Akhil Arora, Robert West, and Martin Gerlach. 2024. Orphan articles: The dark matter of wikipedia. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, pages 100–112.
- Ekaterina Artemova and Barbara Plank. 2023. [Low-resource bilingual dialect lexicon induction with large language models](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 371–385, Tórshavn, Faroe Islands. University of Tartu Library.
- Mikel Artetxe, Vedanuj Goswami, Shruti Bhosale, Angela Fan, and Luke Zettlemoyer. 2023. [Revisiting machine translation for cross-lingual classification](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6489–6499, Singapore. Association for Computational Linguistics.
- Akari Asai, Jungo Kasai, Jonathan Clark, Kenton Lee, Eunsol Choi, and Hannaneh Hajishirzi. 2021. [XOR QA: Cross-lingual open-retrieval question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 547–564, Online. Association for Computational Linguistics.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. 2016. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*.
- Stephen Barbour and Patrick Stevenson. 1998. *Variation im Deutschen: Soziolinguistische Perspektiven*. Walter de Gruyter.
- Adam Berger, Rich Caruana, David Cohn, Dayne Freitag, and Vibhu Mittal. 2000. Bridging the lexical chasm: statistical approaches to answer-finding. In *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 192–199.
- Luiz Bonifacio, Vitor Jeronymo, Hugo Queiroz Abonizio, Israel Campiotti, Marzieh Fadaee, Roberto Lotufo, and Rodrigo Nogueira. 2021. mmarco: A multilingual version of the ms marco passage ranking dataset. *arXiv preprint arXiv:2108.13897*.
- Martin Braschler. 2004. Clef 2003 – overview of results. In *Comparative Evaluation of Multilingual Information Access Systems*, pages 44–63, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Manuel Burghardt, Daniel Granvogl, and Christian Wolff. 2016. [Creating a lexicon of Bavarian dialect by means of Facebook language data and crowdsourcing](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 2029–2033, Portorož, Slovenia. European Language Resources Association (ELRA).
- Andreas Chari, Sean MacAvaney, and Iadh Ounis. 2023. On the effects of regional spelling conventions in retrieval models. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2220–2224.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

- Stephan Elspaß and Robert Möller. 2003. [Atlas zur deutschen Alltagssprache \(AdA\)](#).
- Fahim Faisal, Orevaoghene Ahia, Aarohi Srivastava, Kabir Ahuja, David Chiang, Yulia Tsvetkov, and Antonios Anastasopoulos. 2024. [DIALECTBENCH: An NLP benchmark for dialects, varieties, and closely-related languages](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14412–14454, Bangkok, Thailand. Association for Computational Linguistics.
- Jibril Frej, Didier Schwab, and Jean-Pierre Chevallet. 2020. [WIKIR: A python toolkit for building a large-scale Wikipedia-based English information retrieval dataset](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1926–1933, Marseille, France. European Language Resources Association.
- E.D. Gutiérrez, Ekaterina Shutova, Patricia Lichtenstein, Gerard de Melo, and Luca Gilardi. 2016. [Detecting cross-cultural differences using a multilingual topic model](#). *Transactions of the Association for Computational Linguistics*, 4:47–60.
- Brent Hecht and Darren Gergle. 2010. The tower of babel meets web 2.0: user-generated content and its applications in a multilingual context. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 291–300.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.
- 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](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Omar Khattab and Matei Zaharia. 2020. Colbert: Efficient and effective passage search via contextualized late interaction over bert. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 39–48.
- Dawn Lawrie, James Mayfield, Douglas W. Oard, and Eugene Yang. 2022. HC4: A new suite of test collections for ad hoc CLIR. In *Proceedings of the 44th European Conference on Information Retrieval (ECIR)*.
- Wing Yan Li, Julie Weeds, and David Weir. 2022. [MuSeCLIR: A multiple senses and cross-lingual information retrieval dataset](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1128–1135, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Constantine Lignos, Daniel Cohen, Yen-Chieh Lien, Pratik Mehta, W. Bruce Croft, and Scott Miller. 2019. [The challenges of optimizing machine translation for low resource cross-language information retrieval](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3497–3502, Hong Kong, China. Association for Computational Linguistics.
- Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. Pyserini: A python toolkit for reproducible information retrieval research with sparse and dense representations. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2356–2362.
- Robert Litschko, Ekaterina Artemova, and Barbara Plank. 2023. [Boosting zero-shot cross-lingual retrieval by training on artificially code-switched data](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3096–3108, Toronto, Canada. Association for Computational Linguistics.
- Robert Litschko, Ivan Vulić, Simone Paolo Ponzetto, and Goran Glavaš. 2022. On cross-lingual retrieval with multilingual text encoders. *Information Retrieval Journal*, 25(2):149–183.
- Sean MacAvaney, Craig Macdonald, and Iadh Ounis. 2022. [Streamlining evaluation with ir-measures](#). In *Advances in Information Retrieval - 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10-14, 2022, Proceedings, Part II*, volume 13186 of *Lecture Notes in Computer Science*, pages 305–310. Springer.
- Sean MacAvaney, Luca Soldaini, and Nazli Goharian. 2020. Teaching a new dog old tricks: Resurrecting multilingual retrieval using zero-shot learning. In *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II 42*, pages 246–254. Springer.

- Sean MacAvaney, Andrew Yates, Arman Cohan, and Nazli Goharian. 2019. Cedr: Contextualized embeddings for document ranking. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pages 1101–1104.
- Robert McMaster and Susanna McMaster. 2002. A history of twentieth-century american academic cartography. *Cartography and Geographic Information Science*, 29(3):305–321.
- Meta contributors. 2024. [Wikipedia article depth](#).
- Alice Millour and Karën Fort. 2019. [Unsupervised data augmentation for less-resourced languages with no standardized spelling](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 776–784, Varna, Bulgaria. INCOMA Ltd.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*.
- Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. 2019. Multi-stage document ranking with bert. *arXiv preprint arXiv:1910.14424*.
- Oduwayo Ogundepo, Xinyu Zhang, Shuo Sun, Kevin Duh, and Jimmy Lin. 2022. [AfriCLIRMatrix: Enabling cross-lingual information retrieval for African languages](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8721–8728, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Uma Roy, Noah Constant, Rami Al-Rfou, Aditya Barua, Aaron Phillips, and Yinfei Yang. 2020. [LAReQA: Language-agnostic answer retrieval from a multilingual pool](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5919–5930, Online. Association for Computational Linguistics.
- Shadi Saleh and Pavel Pecina. 2020. [Document translation vs. query translation for cross-lingual information retrieval in the medical domain](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6849–6860, Online. Association for Computational Linguistics.
- Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. [ColBERTv2: Effective and efficient retrieval via lightweight late interaction](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3715–3734, Seattle, United States. Association for Computational Linguistics.
- Shota Sasaki, Shuo Sun, Shigehiko Schamoni, Kevin Duh, and Kentaro Inui. 2018. [Cross-lingual learning-to-rank with shared representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 458–463, New Orleans, Louisiana. Association for Computational Linguistics.
- Shigehiko Schamoni, Felix Hieber, Artem Sokolov, and Stefan Riezler. 2014. [Learning translational and knowledge-based similarities from relevance rankings for cross-language retrieval](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 488–494, Baltimore, Maryland. Association for Computational Linguistics.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2021. [WikiMatrix: Mining 135M parallel sentences in 1620 language pairs from Wikipedia](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1351–1361, Online. Association for Computational Linguistics.
- Lütfi Kerem Senel, Benedikt Ebing, Konul Baghirova, Hinrich Schuetze, and Goran Glavaš. 2024. [Kardeş-NLU: Transfer to low-resource languages with the help of a high-resource cousin – a benchmark and evaluation for Turkic languages](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1672–1688, St. Julian’s, Malta. Association for Computational Linguistics.
- Peng Shi, He Bai, and Jimmy Lin. 2020. [Cross-lingual training of neural models for document ranking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2768–2773, Online. Association for Computational Linguistics.
- Shuo Sun and Kevin Duh. 2020. [CLIRMatrix: A massively large collection of bilingual and multilingual datasets for cross-lingual information retrieval](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4160–4170, Online. Association for Computational Linguistics.
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. [Is ChatGPT good at search? investigating large language models as re-ranking agents](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14918–14937, Singapore. Association for Computational Linguistics.
- Kushal Tatariya, Artur Kulmizev, Wessel Poelman, Esther Ploeger, Marcel Bollmann, Johannes Bjerva, Jiaming Luo, Heather Lent, and Miryam de Lhoneux. 2024. [How good is your Wikipedia?](#) *Preprint*, arXiv:2411.05527.

- Francisco Valentini, Viviana Cotik, Damián Furman, Ivan Bercovich, Edgar Altszyler, and Juan Manuel Pérez. 2024. Messirve: A large-scale spanish information retrieval dataset. *arXiv preprint arXiv:2409.05994*.
- Jannis Vamvas, Noëmi Aeppli, and Rico Sennrich. 2024. [Modular adaptation of multilingual encoders to written Swiss German dialect](#). In *Proceedings of the 1st Workshop on Modular and Open Multilingual NLP (MOOMIN 2024)*, pages 16–23, St Julians, Malta. Association for Computational Linguistics.
- Alastair G. H. Walker and Ommo Wilts. 2001. [Die nordfriesischen Mundarten](#). In Horst Haider Munske, Nils Árhammar, Volker F. Faltings, Jarich F. Hoekstra, Oebele Vries, Alastair G.H. Walker, and Ommo Wilts, editors, *Handbook of Frisian Studies*, pages 284–304. Max Niemeyer Verlag, Berlin, Boston.
- Peter Wiesinger. 1983. [Die Einteilung der deutschen Dialekte](#). In Werner Besch, Ulrich Knoop, Wolfgang Putschke, and Herbert Ernst Wiegand, editors, *Dialektologie: Ein Handbuch zur deutschen und allgemeinen Dialektforschung*, pages 807–900. Walter de Gruyter.
- Mike Zhang and Antonio Toral. 2019. [The effect of translationese in machine translation test sets](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 73–81, Florence, Italy. Association for Computational Linguistics.
- Xinyu Zhang, Sebastian Hofstätter, Patrick Lewis, Raphael Tang, and Jimmy Lin. 2023a. Rank-without-gpt: Building gpt-independent listwise rerankers on open-source large language models. *arXiv preprint arXiv:2312.02969*.
- Xinyu Zhang, Xueguang Ma, Peng Shi, and Jimmy Lin. 2021. [Mr. TyDi: A multi-lingual benchmark for dense retrieval](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 127–137, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xinyu Zhang, Nandan Thakur, Odunayo Ogundepo, Ehsan Kamalloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Mehdi Rezagholizadeh, and Jimmy Lin. 2023b. [MIRACL: A multilingual retrieval dataset covering 18 diverse languages](#). *Transactions of the Association for Computational Linguistics*, 11:1114–1131.
- Wei Zhao, Goran Glavaš, Maxime Peyrard, Yang Gao, Robert West, and Steffen Eger. 2020. [On the limitations of cross-lingual encoders as exposed by reference-free machine translation evaluation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1656–1671, Online. Association for Computational Linguistics.

A Results for Translated Documents

Model	Variations	<i>als</i>	<i>bar</i>	<i>nds</i>	<i>pfl</i>	<i>ksh</i>	avg
BM25	✗	0.598	0.672	0.654	0.602	0.613	0.628
	✓	0.616	0.677	0.651	0.622	0.643	0.642
MonoBERT	✗	0.563	0.491	0.598	0.583	0.538	0.554
	✓	0.586	0.518	0.605	0.598	0.582	0.578
Co1BERT	✗	0.619	0.597	0.554	0.588	0.519	0.575
	✓	0.631	0.597	0.560	0.553	0.539	0.576

Table 6: Cross-dialect retrieval results on **translated documents** of WikiDIR in terms of nDCG@10. We evaluate BM25 and the fine-tuned versions of MonoBERT and ColBERT on two versions of the WikiDIR analysis split. The first version disregards dialect variations (✗) and considers any document relevant that contains a the query. The second version considers any document as relevant if it contains the query or any of its dialect variations (✓).

Document Language	Variations	<i>als</i>	<i>bar</i>	<i>nds</i>	<i>pfl</i>	<i>ksh</i>	avg
Dialect (original)	✗	40.7%	34.7%	24.7%	33.6%	42.3%	35.2%
	✓	25.8%	29.7%	21.3%	21.6%	28.5%	25.4%
German (translated)	✗	69.7%	85.8%	76.9%	79.4%	73.8%	77.1%
	✓	70.5%	80.3%	74.0%	64.6%	69.2%	71.7%

Table 7: Percentage of documents containing an exact match of the query (analysis split).

B Examples of Errors on Translated Documents

Category	German Query (id)	Error Description
Dialect-Specific Variation	<i>Wort</i> (9695)	The query for the German term "Wort" (eng.: word) generates many false positives as "Wort" occurs in many documents. In the Riparian (ksh) reference document, however, the term <i>Woot</i> is not translated correctly.
	<i>Rosine</i> (98086)	The Bavarian dialect term for "Rosine" (eng.: raisin), <i>Weinberl</i> , is not translated at all.
	<i>Weinbau</i> (49763)	The Alemannic term for "Weinbau" (eng.: winegrowing), <i>Räbbau</i> , is referred to with "Rebbaubegriff".
Bag-of-Words	<i>Schlacht bei Kappel</i> (3766713)	Retrieved <i>als</i> documents contain many false positives with the subwords "Schlacht bei", but they miss the proper name "Kappel"
	<i>Alter Friedhof Speyer</i> (4213010)	Retrieved <i>pfl</i> documents contain false positives related to the town name "Speyer", but unrelated to "Alter Friedhof" (eng.: old cemetery).
	<i>Bildende Kunst</i> (714)	The Bavarian term for "Bildende Kunst" (eng.: fine arts), <i>Buidnde Kunst</i> , does not appear in translation. The results for this query contain documents with the term "Kunst" (eng.: art), but are not related to fine arts.

Table 8: Different categories of translation errors in the documents, contributing to weak retrieval performance.