




Enabling Low-Resource Language Retrieval: Establishing Baselines for Urdu MS MARCO

Umer Butt^{1,2} , Stalin Veranasi^{1,2} , and Günter Neumann^{1,2} 

University of Saarland, Germany
German Research Center for Artificial Intelligence (DFKI), Germany
{mubu01, stalin.veranasi, guenter.neumann}@dfki.de

Abstract. As the Information Retrieval (IR) field increasingly recognizes the importance of inclusivity, addressing the needs of low-resource languages remains a significant challenge. This paper introduces the first large-scale Urdu IR dataset, created by translating the MS MARCO dataset through machine translation. We establish baseline results through zero-shot learning for IR in Urdu and subsequently apply the mMARCO multilingual IR methodology to this newly translated dataset. Our findings demonstrate that the fine-tuned model (Urdu-mT5-mMARCO) achieves a Mean Reciprocal Rank (MRR@10) of 0.247 and a Recall@10 of 0.439, representing significant improvements over zero-shot results and showing the potential for expanding IR access for Urdu speakers. By bridging access gaps for speakers of low-resource languages, this work not only advances multilingual IR research but also emphasizes the ethical and societal importance of inclusive IR technologies. This work provides valuable insights into the challenges and solutions for improving language representation and lays the groundwork for future research, especially in South Asian languages, which can benefit from the adaptable methods used in this study.

Keywords: Natural Language Processing · Information Retrieval · Low-resource language.

1 Introduction

Advancements in Natural Language Processing (NLP), particularly transformer architectures, have revolutionized Information Retrieval (IR), enhancing information access worldwide. However, this progress predominantly benefits high-resource languages, leaving low-resource languages like Urdu underrepresented. Urdu is spoken by over 70 million people primarily in South Asia, and it features unique linguistic characteristics, such as its Perso-Arabic script with right-to-left writing direction and rich morphology [2], which pose significant challenges for IR systems. These include complex tokenization, script handling, and morphological analysis, which are also present in some high-resource languages such as Arabic, Chinese, and Russian but have benefited from extensive research and resource availability, making these challenges less of a barrier in those languages.

One of the primary challenges in expanding IR capabilities for low-resource languages like Urdu is the scarcity of large-scale, high-quality datasets necessary for training and evaluating models. Creating such datasets manually is resource-intensive and often impractical for languages with limited digital presence. Consequently, to close this gap, researchers have explored alternative methods, such as machine translation of existing datasets, but it introduces other challenges like translation errors and context loss, adversely affecting IR performance. To our knowledge, no prior work has specifically addressed these issues for IR in Urdu, leaving a gap that our work aims to fill.

In this paper, we address this issue by translating a large-scale MS MARCO passage ranking dataset [3] into Urdu using the IndicTrans2 model [5]. This translation provides a valuable resource for Urdu IR development. We first established baseline performance on our newly created dataset using BM25 and mMarco (a multilingual mT5 model trained on the MS MARCO dataset)[4] in a zero-shot setting, as it was not originally trained on Urdu. We then fine-tune the model on our Urdu-translated dataset. Our experiments demonstrate significant improvements in Mean Reciprocal Rank (MRR) and Recall metrics compared to the zero-shot results, highlighting the potential of our approach.

Our key contributions are:

- Creation of the Urdu MS MARCO Dataset: We provide the first large-scale IR dataset for Urdu by translating the MS MARCO dataset using the IndicTrans2 model, addressing the data scarcity issue.
- Establishment of Baseline IR Performance: We conduct experiments using BM25 and a fine-tuned mT5 model to establish baseline retrieval performance on the Urdu dataset.
- Insights into Low-Resource Language IR: We analyze the challenges and potential solutions for improving IR in Urdu, providing a foundation for future research in similar languages.
- Publicly Available Resources: To facilitate future research in Urdu IR, we make our translated dataset, fine-tuned mT5 model, and code publicly available on Hugging Face ¹ ² and GitHub ³.

2 Related Work

The lack of large-scale, high-quality datasets for low-resource languages has been a significant barrier to advancing IR capabilities. To address this data scarcity, several efforts have created multilingual datasets through machine translation. For example, mMarco[4] extended the MS MARCO dataset[3] to 13 languages using Helsinki NLP models [10] and Google Translate. However, Urdu was not included, partly due to the lack of high-quality translation models for the language at that time.

¹ <https://huggingface.co/datasets/Mavkif/urdu-msmarco-dataset>

² <https://huggingface.co/Mavkif/urdu-mt5-mmarco>

³ https://github.com/UmerTariq1/Urdu_MsMarco_Translation_Retrieval

Recent developments in machine translation, such as the IndicTrans2 model[5], have improved translation quality for South Asian languages, including Urdu. IndicTrans2 achieves a chrF++⁴ score of 68.2 for English-to-Urdu translation, surpassing models like Google Translate. This improvement opens new possibilities for creating large-scale Urdu datasets suitable for training IR models.

In the domain of multilingual NLP, models like mT5[11] have demonstrated strong performance across multiple languages by pre-training on a massive multilingual corpus. Similarly, several works have focused on multilingual IR. ColBERT-X[8] extends the ColBERT[6] architecture to multilingual settings, showing effectiveness in languages included in its training data. Benchmarks like MIRACL[12] and initiatives like CIRAL[1] aim to standardize evaluation and improve cross-lingual retrieval, particularly for low-resource scenarios. MIRACL addresses this by providing a diverse set of 18 languages, including both high and low-resource languages from 10 different language families, with a focus on monolingual retrieval. It features a variety of tasks, including question answering and document retrieval, and incorporates challenges like code-switching and dialectal variations. While MIRACL includes a diverse set of languages, it currently lacks a dedicated Urdu corpus.

Our approach specifically addresses this gap for Urdu, leveraging the improved translation capabilities of IndicTrans2 to create a large-scale Urdu IR dataset. By fine-tuning the mT5 model on this dataset, we address the challenges posed by Urdu’s script and morphology. This work provides a foundation for future research in low-resource languages and contributes to the inclusivity of IR technologies.

3 Dataset Creation

3.1 Ms-Marco dataset

The MS MARCO Passage ranking dataset [3]⁵ is a cornerstone in IR research, comprising over 8.8 million passages and more than 500 thousand unique queries each paired with at least one relevant passage. The commonly used development set includes 6,980 queries, while the training set is organized into 39 million triples, each consisting of a query, a relevant passage, and a non-relevant passage. An example of it can be seen in Table 1. This dataset’s scale and richness have significantly advanced English language IR models. However, replicating such progress for Urdu is challenging due to the absence of equivalent large-scale datasets or the prohibitive costs associated with translating such extensive corpora⁶.

⁴ Character-level F-score, used for capturing partial matches and morphological variations across languages

⁵ <https://microsoft.github.io/msmarco/Datasets.html>

⁶ estimated around \$30,000 for MS Marco, using services like Azure Translator or Google Translation API with their standard rates of 10-20\$ per million characters

3.2 Translation Process

To create the Urdu MS MARCO dataset, we translated the entire English MS MARCO dataset, including both passages and queries from the training and development sets, using the open-source IndicTrans2[5] model. We utilized the distilled 200-million parameter version of IndicTrans2 to translate the entire MS MARCO dataset into Urdu. The process involved tokenizing the text with the IndicTransTokenizer, translating in batches to optimize computational efficiency and leveraging GPU acceleration. Key parameters included a batch size of 32 and a maximum sequence length of 512 tokens to handle lengthy passages. We also performed preprocessing steps, such as normalization and error correction to enhance the quality and usability of the dataset and the model. Overall the translation took around 120 hours to translate the dataset on a single V100 GPU with 32GB Vram.

Table 1. Example of a query and its relevant and non-relevant passage in both English and machine-translated Urdu. More examples can be seen on huggingface dataset webpage

English	Urdu Translation (اردو)
Query : what fruit is native to australia?	Query : آسٹریلیا کا کون سا پھل مقامی ہے؟
Relevant Passage : Passiflora herbertiana. A rare passion fruit native to Australia. Fruits are green-skinned, white fleshed, with an unknown edible rating. Some sources list the fruit as edible, sweet and tasty, while others list the fruits as being bitter and inedible.assiflora herbertiana.	Relevant Passage : پاسفلاورا ہربرتیانا۔ آسٹریلیا کا ایک نایاب پھل ہے۔ پھل سبز جلد والے سفید گوشت والے ہوتے ہیں جن کی خوردنی درجہ بندی نامعلوم ہوتی ہے۔ کچھ ذرائع پھل کو خوردنی مینھا اور لذیذ کے طور پر درج کرتے ہیں جبکہ دوسرے پھلوں کو تلخ اور ہربرتیانا کے طور پر درج کرتے ہیں۔
Non-Relevant Passage: The kola nut is the fruit of the kola tree, a genus (Cola) of trees that are native to the tropical rainforests of Africa.	Non- Relevant Passage: کولا نٹ کولا درخت کا پھل ہے جو درختوں کی ایک جینس کولا ہے جو افریقہ کے اشنکٹبندیی برساتی جنگلات سے تعلق رکھتی۔

4 Experimental Setup

4.1 Baseline Models

Given the absence of prior work on Urdu IR, establishing strong baseline performances is essential for meaningful comparison. We employed BM25, a well-established non-neural retrieval method [9], as our initial retrieval model to retrieve the top 1,000 passages per query[7]. BM25 serves as a fundamental baseline due to its simplicity and effectiveness in various IR tasks. However, its performance tends to be limited for low-resource languages, as it does not account for complex morphological patterns or linguistic nuances present in a language.

To complement BM25, we also evaluated the zero-shot performance of the mMARCO model [4] as a re-ranker. Although Urdu is not included in mMARCO’s

training data, we aimed to assess whether the model could offer effective re-ranking for Urdu text based solely on its multilingual training. This evaluation provides an initial benchmark for the model’s performance on low-resource languages like Urdu.

4.2 Urdu Re-ranker Fine-tuned Model

To enhance retrieval performance beyond the baseline, we fine-tuned the mMARCO model, a fine-tuned version of mT5, on our Urdu-translated MS MARCO dataset, following the original mMARCO training setup for consistency. We set the batch size to 32 with gradient accumulation over 4 steps, resulting in effective batch size of 128 to balance computational efficiency with model performance. The model used a constant learning rate of 0.001 and a dropout rate of 0.1 across all layers, while other hyperparameters remained aligned with the mMARCO setup.

In fine-tuning, we adapted our model for binary relevance classification, generating the urdu word for "yes" for relevant passages and the urdu word for "no" for non-relevant ones, thus framing relevance as a sequence-to-sequence task. This approach provides interpretable relevance decisions, with each query-passage pair yielding a clear binary response.

At inference time, we compute relevance scores by applying a softmax over the logits of the word "yes" and "no" tokens (in urdu script), ranking documents based on the probability of the urdu "yes" token. We fine-tuned the model using the full Urdu-translated MS MARCO dataset, consisting of 39 million triples (query, relevant, and non-relevant passages), with training taking approximately 75 hours on a single A100 GPU with 40GB Vram. This setup allowed the model to adapt effectively and make relevance predictions suited to its linguistic characteristics.

4.3 Evaluation Metrics

To assess the effectiveness of our models, we used common evaluation metrics in information retrieval research to ensure comparability with other studies. These metrics include Mean Reciprocal Rank at 10 (MRR@10), which evaluates the ranking quality by considering the position of the first relevant result within the top 10 retrieved passages. We also used Recall at 10 (Recall@10) to measure the proportion of relevant passages retrieved in the top 10 results, reflecting the system’s ability to find as many relevant passages as possible. These two metrics were also used in the original mMarco benchmark. Additionally, we included Mean Average Precision at 10 (MAP@10) to assess the overall ranking quality across all relevant passages for each query. Finally, we used Normalized Discounted Cumulative Gain at 10 (NDCG@10), a more sophisticated metric that takes into account the position of all relevant passages in the ranking and gives higher weight to those appearing earlier in the list. By employing these metrics, we aimed to provide a comprehensive evaluation of retrieval performance, encompassing various aspects of ranking quality and relevance.

Table 2. Performance Comparison of Retrieval Models on English and Urdu MS-Marco Datasets

Model	Dataset Language	Recall@10	MRR@10	MAP@10	NDCG@10
BM25 (k=1000)	English	0.391	0.187	0.104	0.275
mMARCO Reranker	English	0.639	0.370	0.245	0.498
BM25 (k=1000)	Urdu	0.247	0.121	0.070	0.177
Zero-Shot mMARCO Reranker	Urdu	0.408	0.204	-	-
Urdu mT5-mMARCO (This Work)	Urdu	0.438	0.248	0.159	0.340

5 Results and Discussion

5.1 Performance Comparison

We evaluated the retrieval performance on the Urdu dataset using BM25 for initial retrieval, followed by different re-ranking strategies (see Table 2). The baseline BM25 model is used to retrieve the top 1000 passages per query and it achieves an MRR@10 of **0.121** on the Urdu dataset. The baseline BM25 model exhibited significantly lower performance on the Urdu dataset compared to the English dataset. This disparity underscores the challenges of applying traditional retrieval techniques to morphologically rich languages like Urdu, where effective tokenization requires more nuanced approaches than those typically used for English. The limitations of standard tokenizers for handling Urdu morphology likely contribute to the lower performance observed with BM25.

Applying the zero-shot mMARCO re-ranker to the BM25 outputs improved performance, despite the model not being trained on Urdu. The mMARCO model generalized to Urdu to some extent, likely due to its multilingual training on languages with similar scripts, such as Arabic. This suggests that the model’s semantic understanding and script familiarity enabled it to enhance the ranking of relevant passages over the BM25 baseline.

Our fine-tuned model, Urdu mT5-mMARCO, yielded the most substantial gains. Fine-tuning on the Urdu dataset allowed the model to adapt to the nuances of the language, leading to improvements across all reported metrics. Notably, Urdu mT5-mMARCO achieved a Recall@10 of **0.438**, surpassing both the BM25 baseline (0.247) and the zero-shot mMARCO (0.408). Furthermore, the fine-tuned model exhibited improvements in MRR@10 (**0.248**), exceeding the zero-shot model (0.204) and significantly outperforming the BM25 baseline (0.121). The gains observed in MAP@10 (**0.159**) and NDCG@10 (**0.340**) further demonstrate the positive impact of fine-tuning. This highlights the effectiveness of language-specific fine-tuning, even with machine-translated training data, in enhancing retrieval performance for low-resource languages. The results demonstrate the potential of this approach to bridge the gap in information access for Urdu speakers.

5.2 Impact of Translation Quality and Limitations

We acknowledge that machine translations, like those from IndicTrans2, may not match the precision of human annotations, and translation errors can misalign query and passage semantics, impacting retrieval precision. For example, in table 1, the word *native* is correctly translated in the query (highlighted in green) but omitted in the relevant passage translation (highlighted in red) by the model. Nonetheless, this work serves as a crucial first step toward enabling Urdu IR research. Furthermore, open-sourcing the translation pipeline, data, and model aims to support future multilingual IR projects and address data limitations in low-resource languages beyond Urdu, fostering broader inclusivity in IR research.

6 Future Work and Conclusion

This paper introduces a foundational approach for enabling Information Retrieval (IR) in low-resource languages by translating the MS MARCO dataset into Urdu and fine-tuning the mMARCO model on this data. The baseline BM25 model for Urdu showed significantly lower performance compared to English, underscoring the need for tailored models in Urdu IR. Our results demonstrate that fine-tuning significantly enhances retrieval performance, highlighting the potential of machine translation as a first step in overcoming data limitations for low-resource languages.

This work contributes to more inclusive IR by making retrieval systems accessible to Urdu speakers, with implications for other South Asian languages. Future work could involve incorporating manual verification to refine translation quality and exploring hybrid methods that combine machine and human translations. To better understand the impact of these methods, a comprehensive analysis of translation quality and its quantitative effect on retrieval performance will be conducted. Expanding this approach to other low-resource languages would further support the development of fair, balanced and multilingual IR technologies.

7 Acknowledgement

This research was conducted at the German Research Center for Artificial Intelligence (DFKI) and the University of Saarland. The authors gratefully acknowledge DFKI for providing the necessary hardware resources and the supportive research environment. This work was also supported by the German Federal Ministry of Education and Research (BMBF) as part of the TRAILS project (01IW24005). The authors thank the BMBF for their commitment to advancing research in Artificial Intelligence.

References

1. Adeyemi, M., Oladipo, A., Zhang, X., Alfonso-Hermelo, D., Rezagholizadeh, M., Chen, B., Lin, J.: Ciral at fire 2023: Cross-lingual information retrieval for african languages. In: Proceedings of the 15th Annual Meeting of the Forum for Information Retrieval Evaluation. pp. 4–6 (2023)
2. Arif, S., Azeemi, A.H., Raza, A.A., Athar, A.: Generalists vs. specialists: Evaluating large language models for urdu. arXiv preprint arXiv:2407.04459 (2024)
3. Bajaj, P., Campos, D., Craswell, N., Deng, L., Gao, J., Liu, X., Majumder, R., McNamara, A., Mitra, B., Nguyen, T., et al.: Ms marco: A human generated machine reading comprehension dataset. arXiv preprint arXiv:1611.09268 (2016)
4. Bonifacio, L., Jeronymo, V., Abonizio, H.Q., Campiotti, I., Fadaee, M., Lotufo, R., Nogueira, R.: mmarco: A multilingual version of the ms marco passage ranking dataset. arXiv preprint arXiv:2108.13897 (2021)
5. Gala, J., Chitale, P.A., AK, R., Gumma, V., Doddapaneni, S., Kumar, A., Nawale, J., Sujatha, A., Puduppully, R., Raghavan, V., et al.: Indictrans2: Towards high-quality and accessible machine translation models for all 22 scheduled indian languages. arXiv preprint arXiv:2305.16307 (2023)
6. Khattab, O., Zaharia, M.: 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. pp. 39–48 (2020)
7. Lu, J., Hall, K., Ma, J., Ni, J.: Hyrr: Hybrid infused reranking for passage retrieval. arXiv preprint arXiv:2212.10528 (2022)
8. Nair, S., Yang, E., Lawrie, D., Duh, K., McNamee, P., Murray, K., Mayfield, J., Oard, D.W.: Transfer learning approaches for building cross-language dense retrieval models. In: European Conference on Information Retrieval. pp. 382–396. Springer (2022)
9. Robertson, S.E., Walker, S., Beaulieu, M., Gatford, M., Payne, A.: Okapi at trec-4. Nist Special Publication Sp pp. 73–96 (1996)
10. Tiedemann, J., Aulamo, M., Bakshandaeva, D., Boggia, M., Grönroos, S.A., Nieminen, T., Raganato, A., Scherrer, Y., Vazquez, R., Virpioja, S.: Democratizing neural machine translation with OPUS-MT. Language Resources and Evaluation (58), 713–755 (2023). <https://doi.org/10.1007/s10579-023-09704-w>
11. Xue, L.: mt5: A massively multilingual pre-trained text-to-text transformer. arXiv preprint arXiv:2010.11934 (2020)
12. Zhang, X., Thakur, N., Ogundepo, O., Kamalloo, E., Alfonso-Hermelo, D., Li, X., Liu, Q., Rezagholizadeh, M., Lin, J.: Miracl: A multilingual retrieval dataset covering 18 diverse languages. Transactions of the Association for Computational Linguistics **11**, 1114–1131 (2023)