

The Overlooked Role of Graded Relevance Thresholds in Multilingual Dense Retrieval

Tomer Wullach and Ori Shapira and Amir DN Cohen

OriginAI

{tomew, oris, amirc}@originai.co

Abstract

Dense retrieval models are typically fine-tuned with contrastive learning objectives that require binary relevance judgments, even though relevance is inherently graded. We analyze how graded relevance scores and the threshold used to convert them into binary labels affect multilingual dense retrieval. Using a multilingual dataset with LLM-annotated relevance scores, we examine monolingual, multilingual mixture, and cross-lingual retrieval scenarios. Our findings show that the optimal threshold varies systematically across languages and tasks, often reflecting differences in resource level. A well-chosen threshold can improve effectiveness, reduce the amount of fine-tuning data required, and mitigate annotation noise, whereas a poorly chosen one can degrade performance. We argue that graded relevance is a valuable but underutilized signal for dense retrieval, and that threshold calibration should be treated as a principled component of the fine-tuning pipeline.

1 Introduction

Dense retrieval models have become the standard approach for information retrieval and are now widely used across multilingual search applications (Wang et al., 2024c; Chen et al., 2024). These models are typically fine-tuned with contrastive objectives that rely on distinguishing positive (“relevant”) from negative (“non-relevant”) query-passage pairs (Gao et al., 2021; Izacard et al., 2022). Although this supervision is constructed as a binary decision, relevance is rarely binary in practice (Vakkari and Sormunen, 2004; Sakai, 2021). Passages can vary substantially in how strongly they address a query, and collapsing this variation into a single cutoff removes information that can play a meaningful role during training.

Whether relevance annotations come from human labelers or LLMs, they naturally reflect degrees of relevance. The question is how to leverage

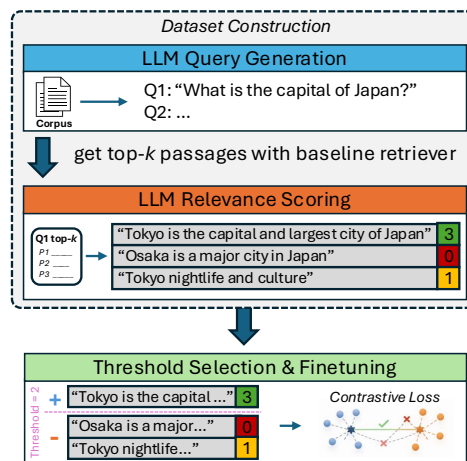


Figure 1: Dense retrieval models are predominantly trained with contrastive learning. Annotating data with graded relevance scores, rather than binary labels, enables selecting an appropriate threshold for positive and negative pairs. This threshold choice has a strong impact on contrastive fine-tuning performance, particularly in multilingual settings.

that graded structure when the model ultimately trains on binary distinctions. Graded labels are converted to positives and negatives by selecting a threshold on the relevance scale. This step determines which examples the model treats as relevant and which it learns to push apart. Even when graded relevance is available, the choice of where to place the threshold is often overlooked as a significant design decision (Scholer and Turpin, 2009; Radlinski and Craswell, 2010). Yet it shapes the diversity of positive examples, the sharpness of the learned decision boundary, and, as our analyses show, even the amount of fine-tuning data needed to obtain improved performance.

In this paper we examine the role of graded relevance and threshold selection in multilingual dense retrieval. We use an LLM-based pipeline to construct a multilingual dataset with fine-grained relevance labels, as illustrated in Figure 1, but our analyses apply equally to datasets annotated by humans

or produced through any multi-level relevance scheme. Across three retrieval scenarios: monolingual retrieval (§4), multilingual mixture fine-tuning (§5), and cross-lingual retrieval (§6), we observe patterns that demonstrate the advantages of configuring the threshold.¹ The optimal threshold varies across languages and tasks, often aligning with factors such as language resource level, pre-training coverage, and the linguistic match between query and corpus. Lower-resource languages tend to benefit from more inclusive thresholds that preserve broader relevance, while higher-resource languages often prefer stricter thresholds during fine-tuning. When languages are mixed during fine-tuning, threshold interactions grow even more pronounced.

Our analysis also highlights the influence of annotation variability (§4.1). We observe disagreement between LLM annotators, especially in lower-resource languages. Thresholding offers a simple way to mitigate the effects of such noise while still making effective use of graded labels.

Taken together, our empirical study demonstrates that while graded relevance provides high-signal data for dense retrieval, its benefits are only fully realized when paired with a thoughtfully chosen threshold. Threshold calibration should therefore be viewed as an important component of the fine-tuning pipeline in multilingual settings, one that can improve effectiveness, stabilize training, and reduce the amount of data required for strong retrieval performance.

2 Background and Related Work

2.1 Dense Retrieval and Contrastive Learning

Dense retrieval has become the dominant paradigm in modern information retrieval due to its ability to semantically represent texts (Zhao et al., 2024). Models under this paradigm learn query and passage embeddings via contrastive learning, that draws relevant pairs near and separates non-relevant ones (Karpukhin et al., 2020; Izacard et al., 2022; Hambarde and Proença, 2023). Although alternative objectives exist, including generative retrieval (Li et al., 2025) and late-interaction models (Khattab and Zaharia, 2020), contrastive learning remains the standard training approach. This motivates our focus on how relevance signals are converted into contrastive supervision.

¹Our resources are available at: <https://github.com/tomerwull/graded-relevance-thresholds>

2.2 Binary versus Graded Relevance Signals

Relevance is inherently graded (Saracevic, 2007), and has long been treated as such in IR *evaluation*, where metrics like nDCG (Järvelin and Kekäläinen, 2002) and ERR (Chapelle et al., 2009) explicitly model degrees of relevance (Roitero et al., 2018). In contrast, *training data* for dense retrieval is predominantly binarized. While prior work has explored the benefits of using more fine-grained relevance judgments for training (Liu et al., 2021b,a, 2022; Tsirigotis et al., 2025), the prevailing practice remains binary labeling, largely for simplicity and scalability.

Several studies point to the limitations of binary relevance labeling. When used for LLM-based annotation, binary prompting can lead to ambiguous or inconsistent judgments because it forces a single decision where gradation is natural (Meng et al., 2025). In contrast, graded labels offer more nuanced supervision (Arabzadeh and Clarke, 2025). Human relevance studies echo these findings. Roitero et al. (2018) show that binary judgments have the lowest agreement with human rankings and emphasize that the appropriate relevance scheme depends on dataset and task requirements.

This body of work highlights that graded relevance is informative and that reducing it to binary form discards useful signals. However, prior work focuses on collecting graded labels rather than understanding how thresholding these labels affects model training, which we address in this paper.

2.3 Multilingual and Cross-Lingual Retrieval

Dense retrieval has also advanced multilingual retrieval, supporting shared training across languages and enabling cross-lingual retrieval (Nair et al., 2022; Wang et al., 2024c). Nonetheless, these models exhibit substantial performance disparities across languages, often correlated with resource levels and pre-training exposure. Benchmarks such as MIRACL (Zhang et al., 2023) and CLIRMatrix (Sun and Duh, 2020) further expose wide variability in multilingual and cross-lingual retrieval quality, shaped by language-specific resource constraints and annotation reliability.

Despite this variability, fine-tuning data in multilingual IR is almost always labeled binarily. This is notable given evidence that multilingual retrieval is especially sensitive to labeling noise. For instance, Shen et al. (2022) show the importance of diverse negative instances, and Huang et al. (2025) find that

Careful negative selection is crucial in low-resource languages. These findings implicitly suggest the value of capturing degrees of relevance, though they do not address threshold calibration directly.

Our analyses extend this work by showing that the optimal relevance threshold depends on language resource level, annotation consistency, and task setting, underscoring the need for flexible threshold calibration in multilingual dense retrieval.

3 Analysis Setup

In the passage retrieval task, a query is provided as input and the most relevant passages from a corpus are returned (Karpukhin et al., 2020; Chang et al., 2020). A dense retrieval model produces vector representations (embeddings) for the query and passages, and matches the query with the closest passages, typically measured via cosine similarity.

Dense retrieval models are trained using contrastive learning objectives, which encourage the model to pull positive (“relevant”) query–passage pairs closer in the embedding space while pushing negatives (“irrelevant”) further apart (Gao et al., 2021; Wang et al., 2024a). Accordingly, in order to fine-tune a model for a particular language or domain, an appropriate set of positive and negative examples is required. Increasingly, practitioners rely on large language models (LLMs) to generate or annotate such data (Wang et al., 2024b; Lee et al., 2024), enabling scalable and cost-effective dataset construction. We follow such an LLM-based pipeline for producing instances used to fine-tune a dense retrieval model, illustrated in Figure 1, which proceeds as follows:

1. **Query generation:** Given a corpus, an LLM is prompted to generate synthetic queries for sampled passages. The generated queries span multiple information-seeking styles, including natural-language questions and keyword-based queries.
2. **Candidate retrieval:** For each generated query, a baseline dense retrieval model retrieves the top- k candidate passages from the corpus (some of which may be irrelevant when there are few results).²

²While candidate selection via baseline retrievers is conventional, potential model-induced bias (see Limitations section) further reinforces the necessity of calibrating thresholds to the specific distribution of candidates.

3. **Relevance labeling:** Each query–passage candidate pair is presented to an LLM, which assigns a relevance score on a predefined scale (possibly binary). This ranking often differs from, and is generally more reliable than, the ranking from the previous step, as also reported by Lee et al. (2024).
4. **Dataset construction:** The resulting fine-tuning dataset contains instances (q, p_i, s_i) , where q is a query, p_i is one of the k retrieved passages, and s_i is the assigned relevance score.

Traditionally, relevance annotation is treated as a binary decision, relevant vs. irrelevant. However, introducing a graded relevance scale can reduce annotation arbitrariness and better reflect the continuum of relevance commonly observed in retrieval tasks (§2). Highly similar query–passage pairs are naturally labeled as clearly relevant, e.g.,

Q: Health effects of excessive and insufficient apoptosis

P: “... Excessive apoptosis causes atrophy, whereas an insufficient amount results in uncontrolled cell proliferation...”

whereas less obvious but still meaningful matches may warrant intermediate scores rather than being forced into a binary “irrelevant”, e.g.,

Q: Knights and pawns

P: “Dunsany’s Chess... One side has standard chess pieces, and the other side has 32 pawns.”

Because contrastive objectives ultimately require binary supervision, a threshold must be applied to convert graded relevance scores into positive and negative instances: scores at or above the threshold become positives, and those below become negatives. Setting this threshold arbitrarily preserves the limitations of purely binary annotation. As the analyses in the subsequent sections show, the threshold choice has consequences.

A higher threshold yields a narrow notion of relevance, exposing the model only to highly similar pairs during training. Lowering the threshold expands the diversity of positive examples, enabling the model to learn more flexible or distant relevance relationships. As we will see, this decision significantly affects performance across languages and tasks.




Figure 2: Number of documents (Wikipedia articles) per language in the MIRACL dataset. In our analyses we use Finnish (fi), Arabic (ar), Japanese (ja), Russian (ru), Spanish (es) and English (en).

3.1 Our Synthetic Dataset

To support our analyses, we prepared a synthetic fine-tuning dataset with graded relevance scores using the pipeline described above.

Source corpus. MIRACL³ (Zhang et al., 2023) is a large corpus covering 18 languages, used for retrieval tasks. To enable analyses across a range of resource levels, we selected Finnish and Arabic as low-resource languages, Japanese as a medium-resource language, Russian and Spanish as high-resource languages, and English as a very high-resource language. Figure 2 shows MIRACL’s distribution of data by language. For each language, we randomly sampled 500K passages outside commonly used development sets. These six corpora form the basis for our synthetic data.⁴

Query generation. We employed GPT-4o⁵ (OpenAI, 2024) to generate language-matched queries designed to resemble web-search intent. A fixed prompt was used across languages, with only the target-language specification modified (see Appendix A).

Baseline dense retrieval model. For candidate retrieval, we used the multilingual-e5-large⁶ embedding model (me5; Wang et al., 2024c), a widely adopted model among practitioners (Enevoldsen et al., 2025), to retrieve a corpus-wide ranking for each query, selecting the top-5

³<https://huggingface.co/datasets/miracl/miracl-corpus>

⁴These six languages were also chosen due to their typological and script diversity, and additionally to align with the test languages available in the cross-lingual evaluation in §6, where the CLIRMatrix MULTI-8 dataset (Sun and Duh, 2020) is used.

⁵We used GPT-4o as it offered a practical balance of multilingual capability and cost for large-scale annotation.

⁶<https://huggingface.co/intfloat/multilingual-e5-large>

passages. The model is pre-trained using MIRACL and related datasets, and we assume its language familiarity roughly follows the distribution shown in Figure 2.

Relevance score annotation. Each query–passage pair was annotated using GPT-4o, assigning a relevance score from 0 to 3, where 3 denotes high relevance and 0 denotes irrelevance. This four-point scale, commonly utilized for model evaluation, has been shown to be effective in this setting (McDonnell et al., 2016; Roitero et al., 2018). The prompt and scale definition appear in Appendix A.

Finally, we sampled down the instances from 500K to 300K, maintaining a similar distribution of scores across the six languages, for fairness in our analyses. Figure 3 (left side) shows the resulting score distributions, and Appendix A describes the downsampling procedure.

With this dataset, we next examine how different threshold choices affect fine-tuning outcomes. Across three use cases, we argue that graded relevance combined with a thoughtfully chosen threshold provides a more flexible and informative supervision signal than binary annotation alone. Adjusting the threshold allows us to explore the trade-off between specificity and coverage in what the model learns as “relevant”. Although our analysis focuses on synthetic data, the considerations around graded relevance and threshold calibration apply equally to any fine-tuning dataset, whether curated by humans or by LLMs. Moreover, the synthetic dataset can be prepared with various component choices, such as different baseline retrieval models or relevance scales. Such variations in the fine-tuning data naturally shift the optimal threshold, necessitating the calibrated approach that we promote.

4 Use-case 1: Monolingual Retrieval

Monolingual passage retrieval involves retrieving passages in the same language as the query and remains the most common retrieval scenario. Using a multilingual dense retrieval model allows holding architecture and training conditions constant across languages while benefiting from shared multilingual representations. This makes monolingual retrieval a natural first setting for analyzing how fine-tuning with graded relevance affects performance across languages, especially given differences in data availability, annotation quality, and pre-training coverage.




Figure 3: Distributions of relevance scores across six languages, as provided by the two LLM annotators (Gpt-4o and Gemma-3-27B). The relevance score ranges from 0 (passage not relevant to query) to 3 (passage very relevant to query). The main difference is apparent in labels 2 and 3.

4.1 Experimental Setup

To evaluate the effect of relevance-aware contrastive fine-tuning, we used our dataset annotated with graded relevance scores for query-passage pairs. From this dataset, we constructed several fine-tuning splits by (i) varying the threshold, $\tau \in \{1, 2, 3\}$, used to designate positive ($s \geq \tau$) and negative ($s < \tau$) pairs, and (ii) varying the number of training instances, from 50K to 300K by intervals of 50K. This allowed us to analyze how both the quantity of supervision and the strictness of the positive/negative boundary influence downstream performance.

We fine-tuned the me5 multilingual dense retriever on each configuration, separately for each of the six languages (technical details in Appendix A). The resulting models were evaluated on the respective language subsets of the MIRACL development set, using nDCG@10 as the primary metric.

To additionally assess the effect of the threshold with respect to the relevance annotator, we repeated the annotation and fine-tuning procedure using relevance scores produced by Gemma-3-27B (Gemma, 2025), a smaller LLM than its counterpart.

GPT vs. Gemma relevance scores. Figure 3 shows the distributional differences between the two LLM annotators. While the score distributions for classes 0 and 1 are broadly similar, the distributions diverge more strongly for the higher relevance labels. GPT-4o assigns label 3 more readily, whereas Gemma-3-27B tends to prefer label 2 over 3. This may indicate that GPT-4o applies a more permissive criterion for identifying strongly relevant passages, or alternatively that Gemma-3-27B is more conservative when making high-confidence relevance judgments.

Figure 4 compares annotator agreement across languages. Agreement is substantially lower for lower-resource languages, mirroring patterns commonly observed in human annotation, where reduced linguistic proficiency leads to noisier judgments. Because LLM abilities vary across languages, particularly for long-tail languages, it is reasonable that relevance scoring becomes less consistent. These differences in annotation consistency further motivate examining how threshold choices affect the preparation of finetuning data. Similarly, in the manner that different LLMs yield varying labels, human annotators often exhibit comparable discrepancies. Regardless of the source, the relevance threshold remains a critical parameter that must be optimized. For this use case we include both LLM annotation sources in our experiments and compare the resulting trends.

4.2 Results

Figure 5 presents the results across languages for varying thresholds and data sizes when using GPT-4o as the annotator (Figure 6 in the appendix shows the corresponding Gemma-3-27B results). We find that increasing the amount of fine-tuning data generally improves monolingual retrieval performance. Initially, small training sets cause a noticeable performance drop relative to the unfine-tuned me5 baseline, suggesting that insufficient supervision can distort rather than enhance the model’s internal representations. As more annotated pairs are used, performance recovers and eventually begins to plateau.⁷

⁷To estimate significance of results, we trained six Finnish models using different random seeds and evaluated all models under the same monolingual setup described in Figure 5. For each threshold and number of training examples, we computed the standard deviation across runs. The observed standard




Figure 4: Agreement heatmaps between the two LLM annotators (GPT-4o and Gemma-3-27B) across languages. Each cell shows the row-normalized frequency of Gemma-3-27B’s scores conditioned on GPT-4o’s score. High-resource languages show stronger diagonal agreement (higher Quadratic Weighted κ), while low-resource languages exhibit greater off-diagonal noise (lower κ). These differences in annotation consistency further motivate examining how threshold choices affect the preparation of finetuning data.

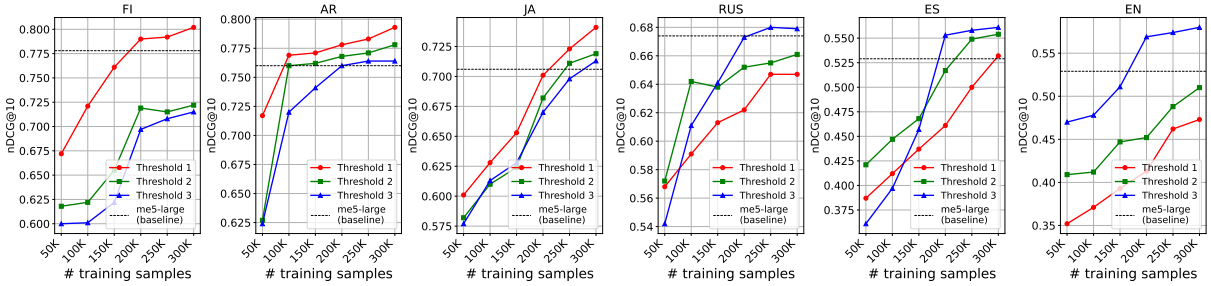


Figure 5: Monolingual retrieval performance (nDCG@10) of me5 on MIRACL (horizontal dashed lines) and after fine-tuning, across six languages. Each subplot shows performance as a function of training set size under different relevance thresholds used to convert graded LLM scores into contrastive training pairs. Lower-resource languages tend to benefit from lower thresholds, whereas higher-resource languages perform better with higher thresholds.

Thresholding strategy plays a critical role. Models trained with poorly chosen relevance thresholds often fail to match the performance of the baseline model. Perhaps more finetuning data would help surpass the baseline. In contrast, well-calibrated thresholds yield clear performance gains with less finetuning data. This underscores that contrastive fine-tuning requires thoughtful binarization rather than naïve thresholding.

A notable pattern emerges across languages. In lower-resource languages, models benefit from lower thresholds, which preserve a broader diversity of relevance relationships. In higher-resource languages, the opposite trend appears, in which stricter thresholds yield better performance. We hypothesize that this arises from differences in the pre-training corpus. For lower-resource languages, the model likely saw fewer examples and a narrower range of positive signals, making additional diversity during fine-tuning beneficial. For higher-resource languages, however, the model likely already encodes rich and diverse relevance distinctions; injecting overly heterogeneous positives may introduce noise and reduce discriminability.

deviation ranges from approximately 3% to 6%, indicating that the results are reasonably stable.

Differences across languages may also reflect properties of the evaluation data. We observe higher nDCG@10 scores for some lower-resource languages, which may indicate that retrieving relevant passages is comparatively easier, perhaps due to more pattern-based query styles or greater lexical overlap between queries and corpus text. Regardless of whether these effects stem from training data, testing data, linguistic differences, or a combination thereof (about which we can only speculate in the scope of this empirical study), the findings collectively highlight the importance of tuning the positive/negative boundary to the target language and domain.

When comparing LLM annotators, we find that models fine-tuned on Gemma-based relevance scores underperform those trained on GPT-based scores. Within the scope of this analysis, we can only speculate about the underlying causes, such as semantic sensitivity or language coverage. Importantly, however, the qualitative trends largely hold across annotators, with only a few exceptions, e.g., Spanish slightly preferring a threshold of 2. While thresholding mitigates some annotation noise, the broader point is that variability in relevance judgments adds yet another consideration when selecting an appro-

Target lang.	Additional lang.	$\tau = 1$	$\tau = 2$	$\tau = 3$
FI	-	0.761	0.655	0.622
	AR	-0.123 (-16.2%)	-0.064 (-9.8%)	-0.100 (-16.1%)
	RUS	+0.049 (+6.43%)	-0.017 (-2.59%)	-0.170 (-27.3%)
	ES	+0.056 (+7.4%)	-0.022 (-3.4%)	-0.052 (-8.4%)
	EN	+0.070 (+9.2%)	+0.000 (+0.0%)	-0.025 (-4.0%)
AR	-	0.771	0.762	0.741
	FI	-0.119 (-15.4%)	-0.185 (-24.3%)	-0.202 (-27.3%)
	RUS	-0.058 (-7.5%)	-0.138 (-18.1%)	-0.141 (-19.0%)
	ES	-0.094 (-12.2%)	-0.155 (-20.3%)	-0.157 (-21.2%)
	EN	-0.049 (-6.35%)	-0.142 (-18.63%)	-0.15 (-20.2%)
RUS	-	0.613	0.638	0.641
	FI	-0.029 (-4.7%)	-0.093 (-14.6%)	-0.084 (-13.1%)
	AR	-0.021 (-3.4%)	-0.11 (-17.2%)	-0.073 (-11.4%)
	ES	+0.012 (+1.95%)	+0.0172 (+2.7%)	+0.01 (+1.5%)
	EN	+0.022 (+3.48%)	+0.019 (+2.97%)	+0.013 (+2.0%)
ES	-	0.437	0.468	0.457
	FI	+0.073 (+16.7%)	+0.036 (+7.7%)	+0.029 (+6.3%)
	AR	+0.054 (+12.4%)	+0.010 (+2.1%)	+0.004 (+0.9%)
	RUS	+0.06 (+13.72%)	+0.041 (+8.7%)	+0.063 (+13.7%)
	EN	+0.064 (+14.6%)	+0.044 (+9.4%)	+0.065 (+14.2%)

Table 1: Monolingual retrieval performance (nDCG@10) on MIRACL when fine-tuning me5 on 150K target-language instances (baseline) and the relative change when adding 150K from another language. The threshold (τ) critically affects how different language mixes influence performance.

appropriate threshold.

Given the consistently higher performance obtained with GPT-based relevance scores, we adopt these scores for the subsequent two use cases.

5 Use-case 2: Finetuning with Mixed Languages

In multilingual dense retrieval, it is often desirable to leverage supervision from other languages, either to compensate for limited labeled data in the target language or to enrich the model’s representational space (e.g., [Thakur et al., 2024](#)). We therefore extend the monolingual passage retrieval setting from the previous section to investigate how the relevance threshold behaves when fine-tuning on a mixture of languages. Specifically, we examine whether the optimal threshold depends solely on the target language or also on the particular combination of languages included in the fine-tuning data.

5.1 Experimental Setup

We adopt the same fine-tuning procedure described in Section 4, but instead of training on a single language, we construct mixed fine-tuning datasets. For each target language, we select 150K instances from that language and shuffle them with an additional 150K instances from another language. A single relevance threshold is applied uniformly across both languages in the mixture. All fine-

tuned models are evaluated on monolingual passage retrieval using the MIRACL development set for the target language, with nDCG@10 as the primary metric.

5.2 Results

Table 1 reports the results for each target language when fine-tuning with 150K instances from that language alone, along with the *relative* performance changes observed when adding 150K instances from another language. Results are shown for thresholds $\tau \in \{1, 2, 3\}$.

We find that multilingual mixtures yield distinct behaviors depending on the target language, the enriching language, and the chosen threshold. Consistent with the monolingual setting, lower-resource target languages benefit from lower thresholds, even when the enrichment language is high-resource. In contrast, higher thresholds lead to increasing performance degradation. For Arabic, adding fine-tuning data from any other language does not improve results and often worsens them, an effect that intensifies as the threshold increases. Russian, in contrast, improves with higher thresholds and benefits comparably from high-resource enrichment languages across all thresholds. However, when the additional language is low-resource, performance drops, though less severely under a lower threshold. Spanish exhibits a more nuanced pattern. When enriched with high-resource languages, gains are modest at threshold 2 and roughly comparable at thresholds 1 and 3. When enriched with Finnish or Arabic (both lower-resourced), however, improvements increase substantially as the threshold decreases, even though Spanish alone achieves stronger baseline performance when the threshold is higher.

These complex (and somewhat confusing) patterns emphasize that in multilingual mixture settings, a single global threshold is clearly suboptimal. The threshold appropriate for the target language may differ from that suitable for the enriching language, and mixing languages can shift the distribution of positive and negative examples in ways that interact complexly with language-specific properties. Overall, the results underscore the importance of carefully selecting relevance thresholds, particularly in multilingual fine-tuning scenarios where both the target language and the composition of the training mixture jointly influence model behavior. We do not theorize here why these patterns emerge, but provide empirical evi-

Queries	Corpus	Low-Med Resource		High Resource		
		AR	JA	RUS	ES	EN
Low-Med Resource	AR	-2.4%	-15.58%	-15.88%	-11.44%	-6.54%
	JA	-17.71%	-4.2%	-11.73%	-11.31%	-6.21%
High Resource	RUS	-14.93%	-10.03%	+3.2%	+8.71%	+3.7%
	ES	-17.71%	-8.36%	+6.02%	+5.1%	+3%
	EN	+0.90%	+5.15%	+2.90%	+2.14%	+7.23%

Table 2: Relative change in cross-lingual retrieval performance when increasing the relevance threshold from 1 to 3 (negative values favor threshold 1, positive values favor threshold 3). Evaluated on the CLIRMatrix MULTI-8 benchmark. Rows correspond to query languages and columns to corpus languages; models are fine-tuned on the query language. Lower-resource languages generally prefer lower thresholds, while pairs of high-resource languages benefit from higher thresholds. See Figure 7 in the appendix for full results.

dence that threshold tuning is essential, particularly when navigating the imbalances inherent in varying language resource levels.

6 Use-case 3: Crosslingual Retrieval

In cross-lingual passage retrieval (e.g., Litschko et al., 2018; Asai et al., 2021), the query is written in one language while relevant passages may appear in another. This scenario is common in multilingual search, where users prefer, or are only able, to formulate queries in their native language even when the corpus is in a different language. Cross-lingual retrieval aims to bridge this linguistic gap by enabling models to retrieve semantically relevant content across languages.

6.1 Experimental Setup

For this use case, we evaluate the same fine-tuned models from the monolingual experiments, specifically those trained with 300K instances, which was the more effective configuration. We test these models on the CLIRMatrix MULTI-8 benchmark (Sun and Duh, 2020), which supports cross-lingual query-passage retrieval. Performance is measured using nDCG@10. The evaluation covers all language pairs for which both query and passage languages appear in the fine-tuning data, excluding Finnish, which is not available in this test set.

6.2 Results

First, our results show that fine-tuning generally improves cross-lingual retrieval performance when the relevance threshold is chosen appropriately. When the threshold is well matched to the language pair, the fine-tuned models consistently outperform the non-fine-tuned baseline; however, an unsuit-

able threshold often leads to underperformance relative to the baseline. These effects are visible in the curves in Figure 7 in the Appendix, which illustrate models’ response to different thresholds across query-corpus combinations. Although more fine-tuning data might eventually mitigate failures caused by a poor threshold, choosing an appropriate threshold avoids most of these cases altogether.

Table 2 summarizes the plots, and quantifies these effects by presenting the relative performance differences when moving from threshold 1 to threshold 3. Negative values, in red, indicate that threshold 1 yields better performance, whereas positive values, in green, indicate that threshold 3 is more beneficial.

Overall, the patterns in Table 2 reveal systematic dependencies between threshold choice and language characteristics. When either the query language or the corpus language is lower-resource, lower thresholds tend to perform better, suggesting that preserving greater variability in relevance pairs is beneficial in such settings. English as a corpus language stands out as a notable exception, where higher thresholds can be advantageous, potentially owing to English’s dominance and extensive representation in pre-training corpora.

When both query *and* corpus languages are high-resource (lower-right region of table), higher thresholds yield stronger results, reflecting already well-calibrated relevance distinctions in these languages.

Despite evaluating on CLIRMatrix rather than MIRACL, the broad trend aligns with the monolingual findings, with English again being the primary exception. These results reinforce that the degree of relevance variability between query and passage (operationalized through the threshold) must be chosen carefully with respect to both the language pair and the retrieval task.

7 Concluding Discussion

Our analysis highlights that graded relevance annotations provide valuable nuance for training multilingual dense retrievers, but their effectiveness depends critically on how they are converted into positive and negative supervision. Because languages and retrieval tasks differ in data availability, semantic variation, and pre-training coverage, the optimal balance between strict and permissive relevance signals is not universal. Threshold calibration is therefore essential since it governs how much variability in relevance the model should internalize,

enabling fine-tuning to adapt appropriately to different linguistic and task conditions.

We also observed notable disagreement between LLM annotators, especially in lower-resource languages. Annotation differences are another cause for controlled thresholding, but it also reinforces the need for mechanisms that can buffer against annotation noise. By controlling how graded labels are mapped to training signals, it can mitigate inconsistencies while still preserving the richer supervision signal that graded relevance offers.

Beyond these methodological considerations, our findings raise important questions about the role of pre-training and evaluation data. The extreme cross-lingual differences in the preferred variability of positive examples, and the unexpectedly high scores in several low-resource languages, indicate that multilingual retrievers may inherit uneven relevance structures shaped by their pre-training environments or by benchmark artifacts. Our analyses shed light on these dynamics, but systematic study of pre-training corpora, annotation behavior, and dataset construction is needed to fully understand and address these discrepancies.

Overall, our findings demonstrate that graded relevance with thoughtful thresholding is a powerful resource for multilingual retrieval. Calibration of the relevance boundary, rather than assuming a one-size-fits-all threshold, should be treated as a central methodological choice in the development of multilingual retrieval systems. Just as hyperparameters are tuned during model training, the relevance threshold should be optimized against a validation set. Our work reveals that such tuning is uniquely critical in multilingual contexts, where optimal values diverge significantly across different languages and tasks.

Limitations

Our analysis covers several representative use cases in multilingual dense retrieval, but many real-world configurations remain unexplored. Retrieval tasks, model architectures, annotation pipelines, datasets, and languages can differ substantially, and we expect threshold behavior to vary accordingly. This diversity is precisely what we aim to highlight. Each setting requires its own assessment rather than relying on a single universal threshold.

Another consideration is that changing the relevance threshold alters the balance between positive and negative training instances. Although the

groups remain large and our results show opposing effects across languages (suggesting that group size is not the dominant factor) this imbalance is inherent to real-world scenarios and should be taken into account when selecting a threshold.

Finally, our dataset construction involves retrieving candidate passages with a baseline model before annotation. Queries that are poorly matched to the corpus may produce more irrelevant candidate passages, introducing a potential bias in the resulting training data. This bias, which arises in any dataset prepared through similar pipelines, could influence fine-tuning behavior and warrants further study.

References

- Negar Arabzadeh and Charles L. A. Clarke. 2025. [Benchmarking llm-based relevance judgment methods](#). In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '25, page 3194–3204, New York, NY, USA. Association for Computing Machinery.
- Akari Asai, Xinyan Yu, Jungo Kasai, and Hanna Hajishirzi. 2021. [One Question Answering Model for Many Languages with Cross-lingual Dense Passage Retrieval](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 7547–7560. Curran Associates, Inc.
- Wei-Cheng Chang, Felix X. Yu, Yin-Wen Chang, Yiming Yang, and Sanjiv Kumar. 2020. [Pre-training Tasks for Embedding-based Large-scale Retrieval](#). In *International Conference on Learning Representations*.
- Olivier Chapelle, Donald Metzler, Ya Zhang, and Pierre Grinspan. 2009. [Expected reciprocal rank for graded relevance](#). In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, CIKM '09, page 621–630, New York, NY, USA. Association for Computing Machinery.
- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. [BGE M3-Embedding: Multi-Lingual, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation](#). *CoRR*, abs/2402.03216.
- Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Marton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Sibli, Dominik Krzeminski, Genta Indra Winata, Saba Sturua, Saiteja Utpala, Mathieu Ciancone, Marion Schaeffer, Diganta Misra, Shreeya Dhakal, Jonathan Rystrom, Roman Solomatin, Omer Veysel Cagatan, and 63 others. 2025. [MMTEB: Massive Multilingual Text Embedding Benchmark](#). In *The Thirteenth International Conference on Learning Representations*.

- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple Contrastive Learning of Sentence Embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Team Gemma. 2025. [Gemma 3 Technical Report](#). Preprint, arXiv:2503.19786.
- Kailash A. Hambarde and Hugo Proença. 2023. [Information Retrieval: Recent Advances and Beyond](#). *IEEE Access*, 11:76581–76604.
- Chao Huang, Fengran Mo, Yufeng Chen, Changhao Guan, Zhenrui Yue, Xinyu Wang, Jinan Xu, and Kaiyu Huang. 2025. [Boosting Data Utilization for Multilingual Dense Retrieval](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 12362–12378, Suzhou, China. Association for Computational Linguistics.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised Dense Information Retrieval with Contrastive Learning](#). *Transactions on Machine Learning Research*.
- Kalervo Järvelin and Jaana Kekäläinen. 2002. [Cumulated gain-based evaluation of IR techniques](#). *ACM Trans. Inf. Syst.*, 20(4):422–446.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense Passage Retrieval for Open-Domain Question Answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, 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, SIGIR '20*, page 39–48, New York, NY, USA. Association for Computing Machinery.
- Jinhyuk Lee, Zhuyun Dai, Xiaoqi Ren, Blair Chen, Daniel Cer, Jeremy R. Cole, Kai Hui, Michael Boratko, Rajvi Kapadia, Wen Ding, Yi Luan, Sai Meher Karthik Duddu, Gustavo Hernandez Abrego, Weiqiang Shi, Nithi Gupta, Aditya Kusupati, Praateek Jain, Siddhartha Reddy Jonnalagadda, Ming-Wei Chang, and Iftexhar Naim. 2024. [Gecko: Versatile Text Embeddings Distilled from Large Language Models](#). Preprint, arXiv:2403.20327.
- Xiaoxi Li, Jiajie Jin, Yujia Zhou, Yuyao Zhang, Peitian Zhang, Yutao Zhu, and Zhicheng Dou. 2025. [From Matching to Generation: A Survey on Generative Information Retrieval](#). *ACM Trans. Inf. Syst.*, 43(3).
- Robert Litschko, Goran Glavaš, Simone Paolo Ponzetto, and Ivan Vulić. 2018. [Unsupervised Cross-Lingual Information Retrieval Using Monolingual Data Only](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18*, page 1253–1256, New York, NY, USA. Association for Computing Machinery.
- Peiyang Liu, Sen Wang, Xi Wang, Wei Ye, and Shikun Zhang. 2021a. [QuadrupletBERT: An Efficient Model For Embedding-Based Large-Scale Retrieval](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3734–3739, Online. Association for Computational Linguistics.
- Peiyang Liu, Xi Wang, Sen Wang, Wei Ye, Xiangyu Xi, and Shikun Zhang. 2021b. [Improving Embedding-based Large-scale Retrieval via Label Enhancement](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 133–142, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Peiyang Liu, Xiangyu Xi, Wei Ye, and Shikun Zhang. 2022. [Label Smoothing for Text Mining](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2210–2219, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Tyler McDonnell, Matthew Lease, Mucahid Kutlu, and Tamer Elsayed. 2016. [Why Is That Relevant? Collecting Annotator Rationales for Relevance Judgments](#). *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 4(1):139–148.
- Chuan Meng, Negar Arabzadeh, Arian Askari, Mohammad Aliannejadi, and Maarten de Rijke. 2025. [Query Performance Prediction Using Relevance Judgments Generated by Large Language Models](#). *ACM Trans. Inf. Syst.*, 43(4).
- Suraj Nair, Eugene Yang, Dawn Lawrie, Kevin Duh, Paul McNamee, Kenton Murray, James Mayfield, and Douglas W. Oard. 2022. [Transfer Learning Approaches for Building Cross-Language Dense Retrieval Models](#). In *Advances in Information Retrieval*, pages 382–396, Cham. Springer International Publishing.
- Team OpenAI. 2024. [GPT-4o System Card](#). Preprint, arXiv:2410.21276.
- Filip Radlinski and Nick Craswell. 2010. [Comparing the sensitivity of information retrieval metrics](#). In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '10*, page 667–674, New York, NY, USA. Association for Computing Machinery. Binarized the 5-point relevance score.

- Kevin Roitero, Eddy Maddalena, Gianluca Demartini, and Stefano Mizzaro. 2018. [On Fine-Grained Relevance Scales](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR '18, page 675–684, New York, NY, USA. Association for Computing Machinery.
- Tetsuya Sakai. 2021. *Graded Relevance*, pages 1–20. Springer Singapore, Singapore.
- Tefko Saracevic. 2007. [Relevance: A review of the literature and a framework for thinking on the notion in information science. Part II: nature and manifestations of relevance](#). *Journal of the American Society for Information Science and Technology*, 58(13):1915–1933.
- Falk Scholer and Andrew Turpin. 2009. [Metric and Relevance Mismatch in Retrieval Evaluation](#). In *Information Retrieval Technology*, pages 50–62, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Tianhao Shen, Mingtong Liu, Ming Zhou, and Deyi Xiong. 2022. [Recovering Gold from Black Sand: Multilingual Dense Passage Retrieval with Hard and False Negative Samples](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10659–10670, Abu Dhabi, United Arab Emirates. 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.
- Nandan Thakur, Jianmo Ni, Gustavo Hernandez Abrego, John Wieting, Jimmy Lin, and Daniel Cer. 2024. [Leveraging LLMs for Synthesizing Training Data Across Many Languages in Multilingual Dense Retrieval](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7699–7724, Mexico City, Mexico. Association for Computational Linguistics.
- Christos Tsirigotis, Vaibhav Adlakha, Joao Monteiro, Aaron Courville, and Perouz Taslakian. 2025. [BiXSE: Improving Dense Retrieval via Probabilistic Graded Relevance Distillation](#). In *Second Conference on Language Modeling*.
- Pertti Vakkari and Eero Sormunen. 2004. [The influence of relevance levels on the effectiveness of interactive information retrieval](#). *Journal of the American Society for Information Science and Technology*, 55(11):963–969.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2024a. [Text Embeddings by Weakly-Supervised Contrastive Pre-training](#). *Preprint*, arXiv:2212.03533.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024b. [Improving Text Embeddings with Large Language Models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11897–11916, Bangkok, Thailand. Association for Computational Linguistics.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024c. [Multilingual E5 Text Embeddings: A Technical Report](#). *Preprint*, arXiv:2402.05672.
- Xinyu Zhang, Nandan Thakur, Odunayo Ogundepo, Ehsan Kamaloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Mehdi Rezagholizadeh, and Jimmy Lin. 2023. [MIRACL: A Multilingual Retrieval Dataset Covering 18 Diverse Languages](#). *Transactions of the Association for Computational Linguistics*, 11:1114–1131.
- Wayne Xin Zhao, Jing Liu, Ruiyang Ren, and Ji-Rong Wen. 2024. [Dense Text Retrieval Based on Pretrained Language Models: A Survey](#). *ACM Trans. Inf. Syst.*, 42(4).

A Implementation Details

LLM annotation prompts. The first phase of preparing the synthetic dataset was generating queries for sampled passages from the MIRACL corpus. We prompted GPT-4o as follows:

```
Your task is to write a <LANG> query
that seeks information existing in
a given context passage.
The query should be relevant to the
passage but does not need to cover
the entire scope of the passage.
The query should be in <LANG>,
short, simple and in the style
of web search query such as
question-based or keyword-based.
```

```
The passage is:
<PASSAGE>
```

An LLM (either GPT-4o or Gemma-3-27B) is later prompted to provide a relevance score for a query-passage pair. The prompt is:

You are a search quality rater evaluating the relevance of passages. Given a query and a passage, you must provide a score on an integer scale of 0 to 3 within brackets [] with the following meanings:

[3] = Highly relevant: The passage is dedicated to the query and contains the exact answer.

[2] = Relevant: The passage provides a partial answer to the query, but it lacks the exact information needed, so it cannot be considered highly relevant.

[1] = Related: The passage seems related to the query but does not answer it.

[0] = Irrelevant: The passage has nothing to do with the query.

Query:

<QUERY>

Passage:

<PASSAGE>

LLM annotation cost. We used the Azure OpenAI API to use GPT-4o, and spent about \$3000 for the query generation phase and \$6000 for the relevance scoring annotations. Annotating with Gemma-3-27B took several hours of GPU server time.

Downsampling the annotated instances. There were 500K query-passage pairs annotated, however to keep an approximately consistent distribution of relevance scores for the six languages, we downsampled to 300K as follows:

1. The language with the smallest amount of non-zero scores (i.e., 1, 2 and 3) was identified.
2. That number of instances was then sampled for the rest of the languages separately.
3. For each language, instances with a score of 0 were sampled until 300K instances total were reached.

As Figure 3 shows, the distribution of scores are quite similar across languages.

Test set sizes. For the languages in our monolingual analyses in use-cases 1 and 2, the MIRACL dev set has between 648 and 2896 instances (queries) depending on the language (see: <https://huggingface.co/datasets/miracl/miracl>). In the CLIRMatrix benchmark, used in the cross-lingual retrieval task (use-case 3), there are 1000 instances (queries) per language.

Fine-tuning the dense retrieval model. We fine-tuned multilingual-e5-large, an encoder-based transformer, using the filtered (q, p) pairs derived from our curated dataset, using the amount of instances that the experiment called for. Models were trained with a batch size of 32 using the InfoNCE loss with in-batch negatives, an initial learning rate of 1×10^{-6} , a 5% warm-up ratio, and for two epochs. To ensure fair comparison across different threshold and data-scale configurations, we kept the training procedure fixed in terms of epochs, optimizer, and learning-rate schedule. While different thresholds lead to different loss behavior, our goal is to compare supervision settings under equal training budgets rather than to individually optimize each configuration.

The models were trained on a server with 24 CPUs, 220 GiB of RAM, and an 80GB NVIDIA A100 GPU. A model took about 5 hours on average to finetune.

Use of AI. We used Copilot with GPT-4o as a backbone for small code improvements when coding the dataset creation pipeline. We also used Copilot to assist in preparing figures and tables for the paper, based on the results of our experiments. Finally, we used ChatGPT 5.1 for assisting in paraphrasing some parts of the text while writing this paper to improve clarity and coherence. Any use of AI was followed by manual verification and editing to maintain correctness and authenticity.

B Supplementing Results

Use-case 1 (monolingual retrieval). Figure 6 shows the results on the monolingual task (use-case 1, §4) when Gemma-3-27B was used as the annotator for relevance scores. This figure is to be compared to Figure 5 in which GPT-4o was used as the annotator instead. As described in Section 4, the performance when using Gemma as the annotator is not as good, but trends are similar with a few differences. This is another example of the importance of setting the threshold sensibly with




Figure 6: Monolingual retrieval performance (nDCG@10) of me5 on MIRACL (horizontal dashed lines) and after fine-tuning, across six languages, when using the relevance scores annotated by Gemma-3-27B. Each subplot shows performance as a function of training set size under different relevance thresholds used to convert graded LLM scores into contrastive training pairs. Lower-resource languages tend to benefit from lower thresholds, whereas higher-resource languages perform better with higher thresholds. See Figure 5 for the results when using GPT-4o as the annotator for relevance scores.

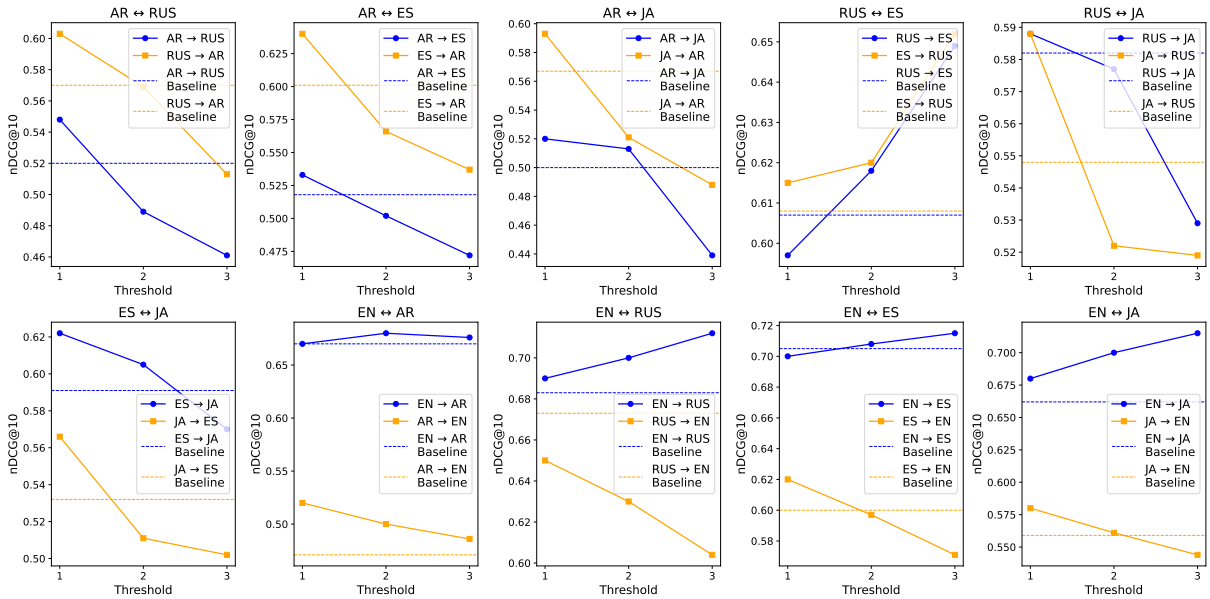


Figure 7: Performance of cross-lingual retrieval for various language pairs across thresholds (represented with $\langle \text{query-lang} \rangle \rightarrow \langle \text{corpus-lang} \rangle$). Each subplot represents a language pair, showing nDCG@10 values for both retrieval directions, evaluated on the CLIRMatrix MULTI-8 benchmark. Dashed horizontal lines indicate the baseline performance achieved using me5 for each direction. Table 2 summarizes the results.

respect to language, data and task.

Use-case 3 (cross-lingual retrieval). Figure 7 presents results for the cross-lingual experiments. Each plot in the figure reports the performance for both retrieval directions ($X \rightarrow Y$ and $Y \rightarrow X$; represented as $\langle \text{query-lang} \rangle \rightarrow \langle \text{corpus-lang} \rangle$), with baseline results from the non-fine-tuned multilingual-e5-large model shown for reference. The figure is also summarized in Table 2. See discussion in Section 6.