

# From MTEB to MTOB: Retrieval-Augmented Classification for Descriptive Grammars

Anonymous ACL submission

## Abstract

Recent advances in language modeling have demonstrated significant improvements in zero-shot capabilities, including in-context learning, instruction following, and machine translation for extremely under-resourced languages (Tanzer et al., 2024). However, many languages with limited written resources rely primarily on descriptions of grammar and vocabulary.

In this paper, we introduce a set of benchmarks to evaluate how well models can extract and classify information from the complex descriptions found in linguistic grammars. We present a Retrieval-Augmented Generation (RAG)-based approach that leverages these descriptions for downstream tasks such as machine translation. Our benchmarks encompass linguistic descriptions for 248 languages across 142 language families, focusing on typological features from WALS (Dryer and Haspelmath, 2013) and Grambank (Skirg rd et al., 2023).

This set of benchmarks offers the first comprehensive evaluation of language models’ in-context ability to accurately interpret and extract linguistic features, providing a critical resource for scaling NLP to low-resource languages. The code and data are publicly available at <https://anonymous.4open.science/r/from-MTEB-to-MTOB>.

## 1 Introduction

The advent of text-based foundational models has accelerated advancements in natural language processing, enhancing multilingual capabilities and applied tasks such as zero-shot machine translation, reading comprehension, and information extraction. Innovations like Machine Translation from One Book (MTOB) (Tanzer et al., 2024) utilize descriptive grammars to improve translation performance for extremely low-resource languages, showing the potential of large-scale language models in bridging linguistic theory and practical NLP applications. Despite the promise, using descriptive grammars

for zero-shot MT presents challenges like terminology variability, non-standard structures, and scattered relevant information.

Additionally, the Massive Text Embedding Benchmark (MTEB) (Muennighoff et al., 2023) provides a thorough evaluation of text embeddings across diverse tasks and languages. A key challenge remains: the effective application of these models to descriptive grammars for languages with scant resources, typically supported by linguistic materials such as grammars and dictionaries.

This paper seeks to address these challenges by providing a systematic framework for extracting information from descriptive grammars and creating a scalable pipeline for descriptive grammar systematization. The key aspect of this approach is Retrieval Augmented Generation (RAG), which allows for the extraction of relevant information from grammars based on a specific typological characteristic (e.g., Order of Subject, Object and Verb). Based on the extracted paragraphs, an LLM determines the value of this characteristic (e.g., Subject-Verb-Object).

In this paper, we present the following contributions:

1. The first scaled linguistic evaluation of the LLM machine reading capabilities on descriptive grammars.
2. A pipeline based on Retrieval Augmented Generation (RAG), which extracts relevant paragraphs from grammars based on a given typological characteristic (for example, WALS 81A: Order of Subject, Object and Verb<sup>1</sup>) and provides them as prompts to an LLM to determine the meaning of these characteristics (for example, Subject-Verb-Object). The pipeline is evaluated through extensive experiments.
3. A benchmark consisting of 700 paragraphs from 14 descriptive grammars, annotated according

<sup>1</sup><https://wals.info/feature/81A>

to whether a linguist can unambiguously determine information about word order (WALS 81A) in the language described, in order to evaluate the quality of information retrieval methods on the task of filtering out the irrelevant paragraphs separately from the RAG pipeline.

4. A benchmark for the RAG pipeline, consisting of 148 grammars for each feature, in order to assess LLMs’ capabilities of determining typological characteristics based on the entire grammar at once and evaluate the effectiveness of different combinations of RAG pipeline components (i.e., information retrieval methods and prompts).

The proposed framework, alongside the presented benchmarks, aims to contribute to the ongoing efforts to improve the quality and efficiency of machine translation systems and to aid linguists in typological research by semi-automating extraction of data from descriptive grammars and. The entire codebase is open-source, licensed under MIT, and available at <https://anonymous.4open.science/r/from-MTEB-to-MTOB>.

## 2 Related Work

Several advances in language modeling outline the possibility to leverage linguistic descriptions: retrieval-augmented generation methods for information extraction and generation, in-context abilities to operate with the descriptive texts in the prompt, and the availability of the texts in the machine-readable format.

### 2.1 Retrieval-Augmented Generation

Retrieval-Augmented Generation (RAG) has become an effective method for augmenting large language models (LLMs) by integrating external retrieval mechanisms. Instead of relying solely on in-model knowledge, RAG enables models to retrieve relevant information from external documents during the generation process (Lewis et al., 2020). This approach has shown promising results in a variety of tasks, such as open-domain question answering and document summarization, by improving the factual accuracy and extending the context of models (Izacard and Grave, 2021). In particular, RAG has been instrumental in improving the performance of LLMs for tasks involving sparse and domain-specific data, where it retrieves external knowledge to complement the model’s inherent capabilities.

RAG’s potential for handling low-resource lan-

guages and complex descriptions, such as linguistic grammars, has not been fully explored. Recent advances in retrieval-based frameworks point toward its applicability in linguistic resource-scarce domains, where available data is often fragmented and incomplete (Gao et al., 2024). This paper leverages the principles of RAG to enhance machine translation and language modeling for under-resourced languages through the use of descriptive grammars.

### 2.2 Zero-shot Learning

Zero-shot learning has become a key area of research in natural language processing (NLP), particularly with the advent of large-scale pre-trained models. The ability of models to generalize to new tasks and languages without explicit task-specific training data is critical for expanding NLP applications to low-resource languages (Brown et al., 2020). LLMs, such as GPT-3 and GPT-4, have demonstrated impressive zero-shot capabilities in tasks ranging from text classification to machine translation, making them an essential tool for under-resourced languages where labeled data is scarce (Raffel et al., 2023).

However, current zero-shot models still struggle with languages that have extremely limited or no monolingual or bilingual corpora. Recent studies, such as (Tanzer et al., 2024) and (Zhang et al., 2024), have shown that integrating linguistic descriptions—such as those found in grammars—can significantly improve zero-shot performance for these languages. In this work, we aim to further this line of research by evaluating the ability of models to utilize descriptive grammars in zero-shot settings.

### 2.3 Grammar Use in NLP

The use of linguistic grammars in NLP has been relatively underexplored, with most efforts focusing on leveraging corpora and parallel data for model training. However, as (Visser, 2022) and others have demonstrated, grammars provide a rich source of structured linguistic knowledge, especially for low-resource languages where corpora are unavailable. According to (Bapna et al., 2022), as cited in (Zhang et al., 2024), 95% of the world’s known languages do not have enough data for LLM training (fewer than 100K sentences), while most languages have linguistic materials available: 60% have a descriptive grammar, and 75% have a dictionary.

Prior work, such as (Erdmann et al., 2017), demonstrates that integrating linguistic knowledge

into a translation model through morphosyntactic parsing can improve translation accuracy for low-resource languages and dialects. However, the challenge of incorporating grammars into NLP models lies in the complexity and formalism of linguistic materials. This paper strives to facilitate further use of grammars in NLP tasks, such as machine translation, through benchmarking the ability of large language models to understand grammars.

## 2.4 Extracting Typological Features from Grammars

Existing research regarding extraction of typological features from grammars precedes the advent of large language models, relying on rule-based methods and earlier developments in classical machine learning and deep learning. The series of works by Virk et al. (Virk et al., 2017); (Virk et al., 2019); (Virk et al., 2020), (Virk et al., 2021) utilize methods that require extensive annotation of semantic frames; (Hammarström et al., 2020) present a method applicable to binary typological features only, and the framework proposed by (Kornilov, 2023) is limited to information retrieval. In this paper, we leverage state-of-the-art language models for the task of typological feature extraction and seek to demonstrate LLMs’ capabilities on the domain of linguistic descriptions.

## 3 Method and Overall Architecture

The basic pipeline for retrieval augmented generation (Naive RAG) (Gao et al., 2024) consists of a database of documents, a retriever, a process of combining the retrieved documents with the prompt, and the LLM generating an answer based on the prompt.

Advanced RAG pipelines described in (Gao et al., 2024) are modifications of different parts of the Naive RAG pipeline. In the context of retrieval augmented generation from a descriptive grammar, the first component of the pipeline—the database of documents—is the grammar itself, hence it is fixed and not as modifiable as for RAG tasks that utilize the Internet or a large database for answering one question. We chose the chunking method that is simple but still context-aware—splitting grammars into paragraphs. We avoid model-based adaptive chunking methods, which would interfere with interpreting the subsequent RAG components due to grammars being a relatively underexplored domain.

The second component of the RAG pipeline is

the retrieval method. We evaluate BM25 (described in (Trotman et al., 2012) and taken as a baseline in (Trotman et al., 2014)), a language-agnostic retriever based on term frequency, and state-of-the-art retrievers/rerankers based on embeddings featured in the Massive Text Embedding Benchmark (MTEB) leaderboard (Muennighoff et al., 2023). Ideally, the chosen embedding-based retrieval methods should respond well to linguistic diversity, since descriptive grammars contain examples in the described language, which may contain diacritic signs and subwords or segments that are rare or unused in English. The weights assigned by the tokenizer to embeddings of such symbols due to their absence in the vocabulary would be random noise, and the resulting embeddings of the paragraphs would have high variance. Therefore, the tokenizer of the chosen retriever should ideally contain byte-level byte pair encoding (BBPE) (Wang et al., 2020).

The third component of the RAG pipeline is the prompt. The prompt format to be used as the baseline only presents the paragraph, the question about the typological characteristic, clarifications regarding what the linguistic term refers to, and a closed set of answers, e.g., for WALS 81A “Dominant Order of Subject, Object and Verb”: “SVO”, “SOV”, “VOS”, “VSO”, “OSV”, “OVS”, “No dominant order”, and additionally “Not enough information” if the dominant word order in the language cannot be inferred from the grammar context. The prompting strategy implemented in the pipeline includes the baseline prompt with additional description of the typological characteristic from WALS or Grambank with examples, as a variation of chain-of-thought prompting (Wei et al., 2022).

The last component of the RAG pipeline is the LLM. We use GPT-4o, OpenAI’s newest flagship model as of May 2024 with increased performance compared to GPT-4 (Achiam et al., 2023): its task is to determine the value of the feature, e.g., “4 cases” for WALS 49A – Number of Cases, based on the prompt and the paragraphs from the descriptive grammar.

In conclusion, the RAG pipeline for descriptive grammars can be called Retrieval Augmented Classification: compared to the more common applications of RAG, the task of the pipeline is to choose one of the values for a linguistic feature from a closed set instead of answering any forms of questions possible, including open-ended ones.

## 4 The Benchmark for Rerankers

### 4.1 Data

Passing the entire grammar to an LLM as a prompt in order to determine the value of a single typological feature is costly and computationally inefficient. Furthermore, the more crucial drawback of passing the unfiltered content to an LLM has been demonstrated in (Shi et al., 2023): quality of LLMs’ responses deteriorates on prompts containing irrelevant context.

Since state-of-the-art retrievers are also LLMs, it would be similarly computationally inefficient to pass the entire grammar to them in order to pass the resulting paragraphs to GPT-4o. Therefore state-of-the-art retrievers become rerankers (one of the advanced additions to the “naive” RAG pipeline): the 50 paragraphs chosen by BM25 are reranked by an LLM retriever, and the resulting top 20 paragraphs are inserted into the prompt for GPT-4o.

The purpose of the benchmark for rerankers is evaluation of state-of-the-art retrievers/rerankers on WALS 81A: Order of Subject, Object and Verb. The benchmark contains 14 grammars written in English: two grammars from each of the six macroareas and two additional grammars for rare word orders (OVS, OSV). Each grammar was split into paragraphs, which were ranked by BM25 using the summary for the English Wikipedia article “Word order” as the query, as proposed in (Kornilov, 2023).

The top 50 paragraphs with the highest ranks assigned by BM25 for each grammar were annotated according to the following principle:

**0** – the paragraph does not mention word order at all;

**1** – the paragraph mentions or describes word order in a construction other than the monotransitive construction (or order of morphemes/phonemes/clitics/etc.);

**2** – the paragraph mentions or describes the word order in the monotransitive construction (in a title of a section, in the table of contents, or in references);

**3** – the paragraph mentions or describes the word order in the monotransitive construction (in a paragraph in the main text);

**4** – the paragraph narrows down the word order in the monotransitive construction to several variants;

**5** – a linguist can unambiguously determine the constituent order in the monotransitive construction

from the paragraph.

Examples of paragraphs for every relevance category are provided in Appendix B.

Among the resulting 700 paragraphs, 38.86% (annotated with 0) are not relevant to order of any elements in a language. Furthermore, 42.43% of the paragraphs (the ones annotated with 1, 2, 3) are potentially misleading data due to describing order of components in a language without explicitly stating the value of WALS 81A. Finally, 18.71% of paragraphs provide relevant information; however, out of the relevant paragraphs, 31.29% are only partially relevant (annotated with 4), and do not sufficiently elaborate on the constituent order in the monotransitive construction in order for the linguist to be able to determine the value of the feature. More detailed data on the benchmark is provided in Appendix C.

In conclusion, this dataset can be used as a benchmark for more advanced information retrieval methods, in order to evaluate their capabilities of filtering out noisy and potentially misleading data.

### 4.2 Results: Evaluating Rerankers

We use the benchmark presented in Section 4.1 to find the best performing retriever/reranker and incorporate it into the RAG pipeline.

As the metric for evaluating the rerankers, we chose NDCG@k (Normalized Discounted Cumulative Gain at k) (Järvelin and Kekäläinen, 2002) over other metrics commonly used for evaluation of information retrieval systems: Recall@k, Mean Average Precision@k (MAP@k), and Mean Reciprocal Rank (MRR), since NDCG@k is the only metric among them that can take into account a scale of more than two relevant ranks: our scale contains six different categories of relevance (0-5) instead of a binary “1 = relevant, 0 = not relevant” distinction.

The rerankers we chose are the 6 models with the best NDCG@10 score on the Massive Text Embedding Benchmark (MTEB) leaderboard (Muenighoff et al., 2023) on the Retrieval task for English as of Aug 29, 2024, along with *SFR-Embedding-Mistral* (Meng et al., 2024), the top 1 model on MTEB as of May 19, 2024. The HuggingFace pages for the models are referenced in Appendix A.

Since all selected embedding models accept instructions, we tested two instruction options - Default Instruct, generic and commonly used for embedding models:



		bge-en-icl	stella_en_1.5B_v5	NV-Retriever-v1	gte-Qwen2-7B-instruct	Linq-Embed-Mistral	SFR-Embedding-2_R	SFR-Embedding-Mistral
Default Instruct	Term Only	0.7387	0.7503	0.6292	0.7598	0.7501	0.7739	0.7381
	Wiki Summary	0.7261	<b>0.7659</b>	<b>0.6693</b>	0.7405	0.7465	0.7624	0.7527
Specific Instruct	Term Only	0.7075	0.7505	0.6492	<b>0.7748</b>	<b>0.7753</b>	<b>0.7750</b>	0.7625
	Wiki Summary	<b>0.7474</b>	0.7601	0.6538	0.7520	0.7690	0.7676	<u><b>0.7776</b></u>

Table 1: NDCG@20 scores on the reranker benchmark. *Term Only* refers to using the sequence “Dominant word order (Order of Subject, Object, and Verb)” as a query, while *Wiki Summary* refers to using the Wikipedia summary from the page “Word order” in English as a query. The best result for each model is shown in **bold**, and the best result across all configuration variants is underlined. The NDCG@20 score for BM25 without a reranker is **0.7494**.

	MTEB Ranking	Grammar Benchmark Ranking	Best Performing Query	Best Performing Instruct	Parameters
bge-en-icl (Xiao et al., 2023)	1	6	Term Only	Default	7.11B
stella_en_1.5B_v5	2	5	Wiki	Default	1.54B
NV-Retriever-v1 (de Souza et al., 2024)	3	7	Wiki	Default	7.11B
gte-Qwen2-7B-instruct (Li et al., 2023)	4	4	Term Only	Specific	7.61B
Linq-Embed-Mistral (Kim et al., 2024)	5	2	Term Only	Specific	7.11B
SFR-Embedding-2_R (Meng* et al., 2024)	6	3	Term Only	Specific	7.11B
SFR-Embedding-Mistral (Meng et al., 2024)	10	1	Wiki	Specific	7.11B

Table 2: Rankings on our benchmark, compared to MTEB, and the best performing configurations for the rerankers. The evaluation of the rerankers on our benchmark requires approximately 1.8 GPU hours in total on 1x NVIDIA A100 40GB.

Given a web search query, retrieve relevant passages that answer the query

and Specific Instruct, tailored to our specific task:

Given a definition of a linguistic feature, retrieve relevant passages that let a linguist unambiguously determine the value of this feature in the described language

Apart from the two variants of the instruct, we use two variants of the query: 1. “Dominant word order (Order of Subject, Object, and Verb)” and 2. the Wikipedia summary from the page “Word order” in English.

In addition to state-of-the-art embedding models, we evaluate BM25 itself as the baseline.

Plots with NDCG@k across all values of k for best performing configurations of each reranker on all grammars are presented in Appendix D, and the plots showcasing comparison between two query types are available in Appendix E.

The results of reranker evaluation are presented in Table 1 and Table 2.

BM25, the baseline retriever, ranks 6th out of 8 and only marginally lags behind the top rerankers in terms of NDCG@20. It appears rational not to take NV-Retriever-V1, which exhibits significantly lower performance compared to other rerankers,

into account, since its sample code on HuggingFace yields different cosine similarity scores with different batch sizes (as of Aug 29, 2024); however, the other reranker outperformed by BM25 on the linguistic domain, *bge-en-icl*, is top 1 on MTEB. Furthermore, the models’ rankings on MTEB and the rankings on our benchmark have a strong negative correlation (Spearman’s  $\rho = -0.8571$ ), indicating that the Information Retrieval subset of MTEB is not suitable for estimating embeddings’ capabilities on the domain of linguistic descriptions.

In conclusion, based on the results on the benchmark described in this section, we have selected *SFR-Embedding-Mistral* with Specific Instruct and Wikipedia Query as the reranker component for the RAG pipeline due to its superior NDCG@20 score compared to other models. Furthermore, the results on the benchmark for rerankers reinforce the decision to select BM25 as the base retriever, since despite being a rule-based method invented decades ago, it has proven itself to be on a similar level compared to state-of-the-art retrievers on the fragment of the linguistic domain presented in our benchmark.

Following the selection of *SFR-Embedding-Mistral* as the reranker, in the following section we proceed to describe the second benchmark created in order to assess the efficacy of the RAG pipeline as a comprehensive system.

## 5 The Benchmark for RAG

### 5.1 Data

The benchmark for the RAG pipeline comprises 148 descriptive grammars. We selected the ostensibly arbitrary number (initially 150) due to benchmarks with fewer than 100 items being unreliable as a tool of assessment: a wrong answer on one item would result in accuracy decreasing by more than one percentage point.

Selecting the grammars in a random fashion would defeat the purpose of the benchmark due to possible biases towards languages having the same descent (a particular family being overrepresented) or being spoken in the same area of the world, without regard to the factual proportions of languages spoken in different areas. Therefore, we used the Genus-Macroarea method described by (Miestamo et al., 2016): in particular, its implementation presented by (Cheveleva, 2023). Identically to the method described in (Cheveleva, 2023), we obtain the proportions of languages across macroareas from the list of genera from WALS, automatically choose descriptive grammars from Glottolog’s (Hammarström et al., 2024) References<sup>2</sup> database, and create the sample anew, placing the limit of one language for each genus. Our sampling strategy differs from the one utilized by (Cheveleva, 2023) in that we limit our sample to grammars written in English, and instead of restricting references to grammar\_sketch and grammar types from Glottolog, we allow references to contain other tags, as long as either grammar\_sketch or grammar is present. The resulting proportions of languages across macroareas are presented in Table 3.

Macroarea	Languages
Africa	29
Australia	9
Eurasia	20
North America	25
Papunesia	39
South America	26
<b>Total</b>	<b>148</b>

Table 3: Languages stratified by macroarea, adapted from (Miestamo et al., 2016) and (Cheveleva, 2023).

We annotated four typological features for the languages described by the grammars in the benchmark. The first feature is identical to the one presented in the benchmark for rerankers: WALS 81A

– Order of Subject, Object, and Verb, as an example of a largely self-explanatory and straightforward feature, which is concentrated in one place in the majority of the grammars: if a paragraph mentions which basic constituent order the language has, the mention is in most cases explicit.

The second annotated feature is from Grambank: GB 107 – “Can standard negation be marked by an affix, clitic or modification of the verb?” Despite being a binary feature, it cannot be reliably extracted by “naive” methods based on term frequency: in the cases when the author of the descriptive literature refers to the negation marker as a marker or a morpheme instead of explicitly calling it a clitic or an affix, the RAG system would have to rely on interlinear glosses and to distinguish between morphological and syntactic phenomena. Furthermore, even in the case when the negation marker is explicitly referred to as a clitic or an affix, it is necessary to determine from context if this marker is phonologically bound solely to the verb (which triggers the feature value = 1) or can be attached to any constituent in the clause (leaving the feature value at 0).

The third feature is a complex composed of seven binary features all related to polar (yes / no) questions; we have chosen it to evaluate the ability of LLMs to reason in the linguistic domain while taking into account several realizations of the same phenomenon simultaneously. This feature will be further referred to as WALS 116A\*: despite being related to WALS 116A, it is more accurately described as an amalgamation of seven separate features from Grambank: GB257, GB260, GB262, GB263, GB264, GB286, and GB291. This feature is essentially a multilabel classification with seven labels: each label/strategy from the set (Interrogative intonation, Interrogative word order, Clause-initial question particle, Clause-medial question particle, Clause-final question particle, Interrogative verb morphology, Tone) is annotated as 1 if it can be used to form polar questions in the described language, and as 0 otherwise.

The last feature is WALS 49A - Number of Cases. It has been chosen due to its quantitative nature, as opposed to binary features determining presence and absence of a particular phenomenon, and its scattered nature: the guideline for WALS 49A takes a liberal approach to determining the number of grammatical cases in a language, allowing adpositional clitics to be considered as case markers; consequently, the relevant information

<sup>2</sup><https://glottolog.org/langdoc>

may be in entirely different sections of the grammar (nominal morphology for "traditional" case markers and syntax for adpositional clitics). Due to Number of Cases being the most time-consuming feature to annotate, we created a separate benchmark for it, consisting partially of grammars from the already existing benchmark and partially of grammars with existing annotations for WALS 49A, maintaining the macroarea proportions and adhering to the one-language-per-genus rule.

The distribution of values for each feature is presented in Appendix H.

## 5.2 Results: Evaluating the RAG Pipeline

In order to evaluate any NLP task on a benchmark, it is crucial to run tests in order to determine if the LLM already possesses knowledge about the feature values in the benchmark. It is different from a contamination test: an LLM may possess knowledge about word order in a particular language from any resource, while a contamination test would determine if the LLM saw the grammar itself during pretraining.

In order to set the baseline for GPT-4o, i.e., to estimate how well it performs on the linguistic domain without additional information from the grammar, we conducted a test on the RAG pipeline excluding the retrieval module. We prompted GPT-4o to determine the values of all benchmark features without the retrieved paragraphs: the only information given to the model consisted of the prompt and the Wikipedia summary for the article about the corresponding feature. The prompts are listed in Appendix F. For each feature, the baseline run was executed ten times in order to capture the variance of the results and to represent presence vs. absence of knowledge more accurately. It is crucial to note that the baseline was evaluated only on those languages where the grammar itself had sufficient information to identify the feature: 136, 146, 121, and 140 values out of 148 for WALS 81A, GB 107, WALS 116A, and WALS 49A respectively.

We subsequently integrated the retriever/reranker component into the pipeline and tested GPT-4o on four prompt configurations: two options for the retrieved information (50 paragraphs from BM25 / 20 paragraphs from a reranker on top of BM25) and two options for prompting (default / with Chain-of-Thought). The Chain-of-Thought prompts are the default RAG prompts concatenated with guidelines and examples from corresponding chapters in WALS and Grambank. We deemed

Grambank chapters particularly suitable for the purpose of Chain-of-Thought prompting, because each chapter comprises: the summary of the feature with the clarifications on ambiguous linguistic terms (i.e., it is explicitly stated in GB263 that only neutral polar questions should be considered in its context, while leading polar questions should be ignored); the step-by-step algorithm intended to instruct human annotators on determining the value of the feature; and examples from the world's languages with interlinear glosses and explanations of the reason why the feature is present (or missing) in the language.

More details regarding the prompt configurations are available in Appendix G.

The temperature value we chose for GPT-4o is 0.2, since our tasks requires the model to incline towards more deterministic behavior.

Running the rerankers on the RAG benchmark for all configurations requires approximately 7.2 GPU hours in total on 1x NVIDIA A100 40GB.

The results for all RAG configurations are presented in Table 4. All RAG configurations outperform the baselines. The observation that macro-averaged F1 scores tend to be higher than micro-averaged F1 scores—which equate to accuracy, given that we do not treat any features as multilabel classifications—suggests that the RAG pipeline is more effective with more frequent classes, struggling to address the class imbalance present in the typological profiles of the world's languages. The results are overall inconsistent, and it is important to note that the Chain-of-Thought approach does not always improve upon its default counterpart, contrary to the expectation for Chain-of-Thought to excel on intensive reasoning tasks.

An additional ablation study is provided in Appendix I.

## 6 Discussion

Expanding the MTOB approach could significantly benefit from standardizing descriptive grammars of various languages into a uniform format, leveraging databases such as Grambank or WALS. However, we can state that in the non-contaminated environment demonstrated in this work the descriptive linguistic texts still pose a significant challenge. While generally machine reading can be perceived as a "solved task", the results on linguistic features show that descriptive grammars remain a non-saturated material showcasing LLMs' weak spots.

	F1 mean	Baseline	BM25	BM25+CoT	Reranker	Rer.+CoT
WALS 81A	micro	0.5551 $\pm$ 0.0359	0.6892	0.7027	0.7095	<b>0.7297</b>
	macro	0.2812 $\pm$ 0.0276	<b>0.7179</b>	0.6097	0.6141	0.6139
	weighted	0.5328 $\pm$ 0.0337	0.6694	0.6890	0.6790	<b>0.6995</b>
GB 107	micro	0.5747 $\pm$ 0.0350	0.6622	0.6554	0.6757	<b>0.7095</b>
	macro	0.5740 $\pm$ 0.0354	0.5718	0.5731	0.6011	<b>0.6546</b>
	weighted	0.5724 $\pm$ 0.0361	0.5957	0.5959	0.6221	<b>0.6713</b>
WALS 49A	micro	0.3986 $\pm$ 0.0238	0.5203	0.5405	<b>0.5541</b>	0.5135
	macro	0.2216 $\pm$ 0.0158	<b>0.4711</b>	0.4494	0.4332	0.4042
	weighted	0.3453 $\pm$ 0.0237	0.5314	0.5542	<b>0.5605</b>	0.5185
interrog. intonation only	micro	0.4504 $\pm$ 0.0396	0.8446	<b>0.8986</b>	0.8378	0.8784
	macro	0.4240 $\pm$ 0.0456	0.8392	<b>0.8946</b>	0.8335	0.8739
	weighted	0.4068 $\pm$ 0.0484	0.8480	<b>0.9007</b>	0.8415	0.8810
interrog. word order	micro	0.9653 $\pm$ 0.0076	<b>0.9932</b>	0.9865	0.9865	0.9797
	macro	0.6641 $\pm$ 0.0460	<b>0.9427</b>	0.8965	0.8715	0.8281
	weighted	0.9611 $\pm$ 0.0064	<b>0.9936</b>	0.9878	0.9865	0.9808
clause- initial particle	micro	0.7264 $\pm$ 0.0402	0.9054	0.9189	<b>0.9459</b>	0.9257
	macro	0.5446 $\pm$ 0.0502	0.7977	0.8335	<b>0.8890</b>	0.8504
	weighted	0.7371 $\pm$ 0.0337	0.9054	0.9205	<b>0.9470</b>	0.9278
clause- final particle	micro	0.4909 $\pm$ 0.0409	0.7230	0.7568	0.7027	<b>0.7838</b>
	macro	0.4808 $\pm$ 0.0415	0.7173	0.7463	0.6972	<b>0.7717</b>
	weighted	0.4661 $\pm$ 0.0428	0.7314	0.7644	0.7116	<b>0.7902</b>
clause- medial particle	micro	0.6760 $\pm$ 0.0367	0.8311	<b>0.8851</b>	0.8378	0.8581
	macro	0.5011 $\pm$ 0.0544	0.7638	<b>0.8201</b>	0.7709	0.7832
	weighted	0.6644 $\pm$ 0.0355	0.8439	<b>0.8888</b>	0.8496	0.8641
interrog. verb morphology	micro	0.7678 $\pm$ 0.0185	0.7973	<b>0.8311</b>	0.8243	<b>0.8311</b>
	macro	0.4533 $\pm$ 0.0311	0.7075	0.7354	0.7282	<b>0.7419</b>
	weighted	0.7102 $\pm$ 0.0160	0.8192	0.8451	0.8396	<b>0.8465</b>
tone	micro	0.9273 $\pm$ 0.0128	0.9257	0.9595	0.9662	<b>0.9730</b>
	macro	0.5147 $\pm$ 0.0573	0.7407	0.8225	<b>0.8594</b>	0.8501
	weighted	0.9104 $\pm$ 0.0121	0.9390	0.9637	0.9704	<b>0.9730</b>

Table 4: F1 scores across all configurations of the pipeline. All five configurations contain prompts from Appendix F with Wikipedia summaries for corresponding features. The *Baseline* column refers to prompting GPT-4o without materials from grammars (i.e., without RAG). *BM25* refers to using 50 paragraphs retrieved by BM25. *Reranker* refers to using 20 paragraphs selected from these 50 by *SFR-Embedding-Mistral*. *CoT* refers to Chain-of-Thought (adding instructions and examples from WALS or Grambank to the prompt). The best F1 value in each row is shown in **bold**. The values after the  $\pm$  sign in the *Baseline* column are the sample standard deviation across 10 runs. Each of the remaining four configurations was executed only once.

## 7 Conclusion

In this paper, we introduced two benchmarks for evaluation of methods that combine Retrieval Augmented Generation with large language models to extract and classify typological features from descriptive grammars.

We also provided an open-source pipeline for linguistic information extraction, which has significant potential to improve NLP applications for under-resourced languages.

The proposed pipeline, alongside the presented benchmarks, revealed that BM25, a language-agnostic information retrieval method, is comparable in quality to state-of-the-art embedding-based methods on the task of retrieving information from descriptive grammars, and can be used as a RAG component on the linguistic domain. Furthermore, despite the notion that the machine reading task has mostly been resolved, the complexities inherent in linguistic texts still present challenges. While the

advancements in language models have made significant strides in handling various types of texts, the results obtained on our benchmarks suggest that it remains premature to assert their effectiveness on the domain of linguistic descriptions.

Our contributions lay the groundwork for extending the capabilities of LLMs to handle complex linguistic data, such as grammatical descriptions. This work represents a crucial step toward better supporting low-resource languages in NLP. Future work could further optimize the retrieval and classification processes, expand the benchmark to include more languages, and explore practical applications of information extraction on the linguistic domain, such as cross-lingual typological analysis and machine translation for extremely low-resource languages.



## 8 Limitations

One of the limitations of our work is the number of languages presented: despite efforts to choose a stratification method that would yield a representative sample, 148 languages (the chosen number for each feature) constitute only about 2% of the world’s languages.

Another limitation is potential presence of erroneous data. WALS and Grambank currently list 2,662 and 2,467 languages respectively, which falls into the range of 30-40% of the world’s known languages (fewer than 60% with grammars available). Furthermore, most language profiles are not full: there are only 5 and 841 languages annotated for the least “popular” feature in WALS and Grambank respectively. Additionally, (Baylor et al., 2023) demonstrate that typological databases contain erroneous entries and discrepancies, reporting 69.04% average agreement between WALS and Grambank across six typological features (agreement score for each feature obtained by averaging scores for macroareas). A portion of the discrepancies is most likely attributable to human error. Despite extensive efforts to eliminate the errors and fill the gaps, our benchmark is not to be considered perfect data either, since the first author was the only annotator.

Furthermore, due to high computational costs of running RAG on GPT-4o with extensive excerpts from grammars, we limited our experiments (excluding the baseline, which was executed 10 times) to running every RAG configuration only once.

Finally, one of the most crucial limitations of this paper is the inability to make the benchmark for the RAG pipeline fully open-source, since the majority of grammars in the benchmark are under copyright.

## 9 Ethical Statement

In pursuing the advancement of machine translation and natural language processing through the use of descriptive grammars, we must address critical ethical concerns, particularly relating to copyright. Most grammars in the benchmark have copyright licenses which prohibit their reproduction fully.

The benchmark for rerankers and the benchmark for RAG differ in the following way: in the benchmark for rerankers, each individual paragraph is annotated according to relevance to WALS 81A (to be used for retrieval assessment), whereas the benchmark for RAG does not have annotations for individual paragraphs. Therefore, we release our

codebase and materials under MIT and publish the first benchmark fully, since it can be considered as a derivative work, and its publication for research purposes falls under fair use.

To ensure compliance with copyright laws and respect for intellectual property rights, we do not publish the second benchmark fully; we only publish the list of grammars and the list of page numbers where the relevant information for each feature can be found.

Furthermore, we advocate for the development of collaborative agreements with copyright holders in order to obtain permission for making at least some of the full texts of the grammars available as open-source. We believe that addressing these copyright issues is crucial for the sustainable and ethical advancement of language technologies.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- Ankur Bapna, Isaac Caswell, Julia Kreutzer, Orhan Firat, Daan van Esch, Aditya Siddhant, Mengmeng Niu, Pallavi Baljekar, Xavier Garcia, Wolfgang Macherey, et al. 2022. [Building machine translation systems for the next thousand languages](#). *arXiv preprint arXiv:2205.03983*.
- Emi Baylor, Esther Ploeger, and Johannes Bjerva. 2023. [The past, present, and future of typological databases in nlp](#). *arXiv preprint arXiv:2310.13440*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.
- Akua Asantewaa Campbell. 2017. [A grammar of Gã](#). Ph.D. thesis.
- Anastasiia Cheveleva. 2023. [Neutralization of gender values in the plural](#). Bachelor’s thesis, HSE University.
- Gabriel de Souza, P. Moreira, Radek Osmulski, Mengyao Xu, Ronay Ak, Benedikt Schifferer, and Even Oldridge. 2024. [Nv-retriever: Improving text](#)

769	embedding models with effective hard-negative min-	Albert Kornilov. 2023. <a href="#">Multilingual automatic extrac-</a>	821
770	ing. <i>Preprint</i> , arXiv:2407.15831.	tion of linguistic data from grammars. In <i>Proceed-</i>	822
771	Matthew S. Dryer and Martin Haspelmath, editors. 2013.	ings of the Second Workshop on NLP Applications to	823
772	<i>WALS Online (v2020.3)</i> . Zenodo.	<i>Field Linguistics</i> , pages 86–94, Dubrovnik, Croatia.	824
773	John Asher Dunn. 1979. <i>A reference grammar for the</i>	Association for Computational Linguistics.	825
774	<i>Coast Tsimshian language</i> . University of Ottawa	Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio	826
775	Press.	Petroni, Vladimir Karpukhin, Naman Goyal, Hein-	827
776	John A Elliott. 2021. <i>A grammar of Enxet Sur</i> . Ph.D.	rich Küttler, Mike Lewis, Wen-tau Yih, Tim Rock-	828
777	thesis, University of Hawai’i at Manoa.	täschel, et al. 2020. <a href="#">Retrieval-augmented generation</a>	829
778	Alexander Erdmann, Nizar Habash, Dima Taji, and	for knowledge-intensive nlp tasks. <i>Advances in Neu-</i>	830
779	Houda Bouamor. 2017. <a href="#">Low resourced machine</a>	<i>ral Information Processing Systems</i> , 33:9459–9474.	831
780	translation via morpho-syntactic modeling: the case	Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long,	832
781	of dialectal arabic. <i>arXiv preprint arXiv:1712.06273</i> .	Pengjun Xie, and Meishan Zhang. 2023. <a href="#">Towards</a>	833
782	Lysbeth Julie Ford. 1998. <i>A description of the Emmi lan-</i>	general text embeddings with multi-stage contrastive	834
783	<i>guage of the Northern Territory of Australia</i> . Ph.D.	learning. <i>arXiv preprint arXiv:2308.03281</i> .	835
784	thesis, Australian Natl. Univ. Canberra.	Rui Meng*, Ye Liu*, Shafiq Rayhan Joty, Caiming	836
785	Diana Forker. 2013. <i>A grammar of Hinuq</i> , volume 63.	Xiong, Yingbo Zhou, and Semih Yavuz. 2024. <a href="#">Sfr-</a>	837
786	Walter de Gruyter.	embedding-2: Advanced text embedding with multi-	838
787	Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia,	stage training.	839
788	Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang,	Rui Meng, Ye Liu, Shafiq Rayhan Joty, Caiming	840
789	and Haofen Wang. 2024. <a href="#">Retrieval-augmented gener-</a>	Xiong, Yingbo Zhou, and Semih Yavuz. 2024. <a href="#">Sfr-</a>	841
790	ation for large language models: A survey. <i>Preprint</i> ,	embedding-mistral:enhance text retrieval with trans-	842
791	arXiv:2312.10997.	fer learning. Salesforce AI Research Blog.	843
792	Stefan Georg. 2007. <i>A descriptive grammar of Ket</i>	Matti Miestamo, Dik Bakker, and Antti Arppe. 2016.	844
793	<i>(Yenisei-Ostyak): part 1: introduction, phonology</i>	<a href="#">Sampling for variety</a> . <i>Linguistic Typology</i> , 20(2):233–	845
794	<i>and morphology</i> , volume 1. Global Oriental.	296.	846
795	Harald Hammarström, One-Soon Her, and Marc	Lawrence Morgan. 1991. <a href="#">A description of the kutenai</a>	847
796	Allasonnière-Tang. 2020. <a href="#">Term spotting: A quick-</a>	<a href="#">language</a> .	848
797	and-dirty method for extracting typological features	Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and	849
798	of language from grammatical descriptions. In <i>Se-</i>	Nils Reimers. 2023. <a href="#">Mteb: Massive text embedding</a>	850
799	<i>lected contributions from the Eighth Swedish Lan-</i>	<a href="#">benchmark</a> . <i>Preprint</i> , arXiv:2210.07316.	851
800	<i>guage Technology Conference (SLTC-2020)</i> , pages	Paul Newman. 2002. <a href="#">The hausa language. an encyclo-</a>	852
801	27–34.	<a href="#">pedic reference grammar</a> . <i>Journal of the American</i>	853
802	Harald Hammarström, Robert Forkel, Martin Haspel-	<i>Oriental Society</i> , 122:97.	854
803	math, and Sebastian Bank, editors. 2024. <a href="#">Glottolog</a>	Knut J. Olawsky. 2006. <i>A Grammar of Urarina</i> . De	855
804	<i>5.0</i> . Zenodo.	Gruyter Mouton, Berlin, New York.	856
805	Birgit Hellwig. 2019. <i>A Grammar of Qaqet</i> . De Gruyter	Colin Raffel, Noam Shazeer, Adam Roberts, Katherine	857
806	Mouton, Berlin, Boston.	Lee, Sharan Narang, Michael Matena, Yanqi Zhou,	858
807	Gautier Izacard and Edouard Grave. 2021. <a href="#">Lever-</a>	Wei Li, and Peter J. Liu. 2023. <a href="#">Exploring the limits</a>	859
808	<a href="#">aging passage retrieval with generative models</a>	<a href="#">of transfer learning with a unified text-to-text trans-</a>	860
809	<a href="#">for open domain question answering</a> . <i>Preprint</i> ,	<a href="#">former</a> . <i>Preprint</i> , arXiv:1910.10683.	861
810	arXiv:2007.01282.	Jeanette Sakel. 2004. <i>A grammar of Mosetén</i> . De	862
811	Kalervo Järvelin and Jaana Kekäläinen. 2002. <a href="#">Cu-</a>	Gruyter Mouton, Berlin, New York.	863
812	<a href="#">mulated gain-based evaluation of ir techniques</a> .	Freda Shi, Xinyun Chen, Kanishka Misra, Nathan	864
813	<i>ACM Transactions on Information Systems (TOIS)</i> ,	Scales, David Dohan, Ed H Chi, Nathanael Schärli,	865
814	20(4):422–446.	and Denny Zhou. 2023. <a href="#">Large language models can</a>	866
815	Junseong Kim, Seolhwa Lee, Jihoon Kwon, Sangmo	<a href="#">be easily distracted by irrelevant context</a> . In <i>Inter-</i>	867
816	Gu, Yejin Kim, Minkyung Cho, Jy yong Sohn, and	<i>national Conference on Machine Learning</i> , pages	868
817	Chanyeol Choi. 2024. <a href="#">Linq-embed-mistral:elevating</a>	31210–31227. PMLR.	869
818	<a href="#">text retrieval with improved gpt data through task-</a>	Hedvig Skirgård, Hannah J. Haynie, Damián E.	870
819	<a href="#">specific control and quality refinement</a> . Linq AI Re-	Blasi, Harald Hammarström, Jeremy Collins, Jay J.	871
820	search Blog.	Latarche, Jakob Lesage, Tobias Weber, Alena	872
		Witzlack-Makarevich, Sam Passmore, Angela Chira,	873

874	Luke Maurits, Russell Dinnage, Michael Dunn,	<i>In Proceedings of the International Conference on</i>	933
875	Ger Reesink, Ruth Singer, Claire Bower, Patience	<i>Recent Advances in Natural Language Processing</i>	934
876	Epps, Jane Hill, Outi Vesakoski, Martine Robbeets,	( <i>RANLP 2021</i> ), pages 1480–1489.	935
877	Noor Karolin Abbas, Daniel Auer, Nancy A. Bakker,		
878	Giulia Barbos, Robert D. Borges, Swintha Danielsen,	Shafqat Mumtaz Virk, Harald Hammarström, Lars	936
879	Luise Dorenbusch, Ella Dorn, John Elliott, Gi-	Borin, Markus Forsberg, and Søren Wichmann. 2020.	937
880	ada Falcone, Jana Fischer, Yustinus Ghanggo Ate,	<a href="#">From linguistic descriptions to language profiles</a> . In	938
881	Hannah Gibson, Hans-Philipp Göbel, Jemima A.	<i>Proceedings of the 7th Workshop on Linked Data in</i>	939
882	Goodall, Victoria Gruner, Andrew Harvey, Re-	<i>Linguistics (LDL-2020)</i> , pages 23–27.	940
883	bekah Hayes, Leonard Heer, Roberto E. Herrera Mi-		
884	randana, Nataliia Hübler, Biu Huntington-Rainey,	Shafqat Mumtaz Virk, Azam Sheikh Muhammad, Lars	941
885	Jessica K. Ivani, Marilen Johns, Erika Just, Eri	Borin, Muhammad Irfan Aslam, Saania Iqbal, and	942
886	Kashima, Carolina Kipf, Janina V. Klingenberg,	Nazia Khurram. 2019. <a href="#">Exploiting frame-semantics</a>	943
887	Nikita König, Aikaterina Koti, Richard G. A. Kowa-	<a href="#">and frame-semantic parsing for automatic extraction</a>	944
888	lik, Olga Krasnoukhova, Nora L.M. Lindvall, Mandy	<a href="#">of typological information from descriptive gram-</a>	945
889	Lorenzen, Hannah Lutzenberger, Tônia R.A. Mar-	<a href="#">mars of natural languages</a> . In <i>Proceedings of the In-</i>	946
890	tins, Celia Mata German, Suzanne van der Meer,	<i>ternational Conference on Recent Advances in Natu-</i>	947
891	Jaime Montoya Samamé, Michael Müller, Saliha	<i>ral Language Processing (RANLP 2019)</i> , pages 1247–	948
892	Muradoğlu, Kelsey Neely, Johanna Nickel, Miina	1256.	949
893	Norvik, Cheryl Akinyi Oluoch, Jesse Peacock, In-		
894	dia O.C. Pearey, Naomi Peck, Stephanie Petit, Sören	Eline Visser. 2022. <a href="#">A grammar of Kalamang</a> . Num-	950
895	Pieper, Mariana Poblete, Daniel Prestipino, Linda	ber 4 in Comprehensive Grammar Library. Language	951
896	Raabe, Amna Raja, Janis Reimringer, Sydney C. Rey,	Science Press, Berlin.	952
897	Julia Rizaew, Eloisa Ruppert, Kim K. Salmon, Jill		
898	Sammet, Rhiannon Schembri, Lars Schlabbach, Fred-	Changhan Wang, Kyunghyun Cho, and Jiatao Gu. 2020.	953
899	erick W.P. Schmidt, Amalia Skilton, Wikaliler Daniel	<a href="#">Neural machine translation with byte-level subwords</a> .	954
900	Smith, Hilário de Sousa, Kristin Sverredal, Daniel	In <i>Proceedings of the AAAI conference on artificial</i>	955
901	Valle, Javier Vera, Judith Voß, Tim Witte, Henry Wu,	<i>intelligence</i> , volume 34, pages 9154–9160.	956
902	Stephanie Yam, Jingting Ye, Maisie Yong, Tessa Yu-		
903	ditha, Roberto Zariquiey, Robert Forkel, Nicholas	Claudia Wegener. 2012. <a href="#">A Grammar of Savosavo</a> . De	957
904	Evans, Stephen C. Levinson, Martin Haspelmath,	Gruyter Mouton, Berlin, Boston.	958
905	Simon J. Greenhill, Quentin D. Atkinson, and Rus-		
906	sell D. Gray. 2023. <a href="#">Grambank reveals global patterns</a>	Jason Wei, Xuezhong Wang, Dale Schuurmans, Maarten	959
907	<a href="#">in the structural diversity of the world’s languages</a> .	Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou,	960
908	<i>Science Advances</i> , 9.	et al. 2022. <a href="#">Chain-of-thought prompting elicits reason-</a>	961
		<a href="#">ing in large language models</a> . <i>Advances in neural</i>	962
		<i>information processing systems</i> , 35:24824–24837.	963
909	Garrett Tanzer, Mirac Suzgun, Eline Visser, Dan Juraf-		
910	sky, and Luke Melas-Kyriazi. 2024. <a href="#">A benchmark</a>	EM Helen Weir. 1986. <a href="#">Footprints of yesterday’s syntax:</a>	964
911	<a href="#">for learning to translate a new language from one</a>	<a href="#">Diachronic development of certain verb prefixes in</a>	965
912	<a href="#">grammar book</a> . <i>Preprint</i> , arXiv:2309.16575.	<a href="#">an osv language (nadëb)</a> . <i>Lingua</i> , 68(4):291–316.	966
913	Andrew Trotman, Xiangfei Jia, and Matt Crane. 2012.		
914	<a href="#">Towards an efficient and effective search engine</a> . In	Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas	967
915	<i>OSIR@ SIGIR</i> , pages 40–47.	Muennighoff. 2023. <a href="#">C-pack: Packaged resources</a>	968
916	Andrew Trotman, Antti Puurula, and Blake Burgess.	<a href="#">to advance general chinese embedding</a> . <i>Preprint</i> ,	969
917	2014. <a href="#">Improvements to bm25 and language models</a>	arXiv:2309.07597.	970
918	<a href="#">examined</a> . In <i>Proceedings of the 19th Australasian</i>		
919	<i>Document Computing Symposium</i> , pages 58–65.	Kexun Zhang, Yee Man Choi, Zhenqiao Song, Taiqi	971
920	Tasaku Tsunoda. 2012. <a href="#">A Grammar of Warrongo</a> . De	He, William Yang Wang, and Lei Li. 2024. <a href="#">Hire</a>	972
921	Gruyter Mouton, Berlin, Boston.	<a href="#">a linguist! Learning endangered languages with</a>	973
922	Shafqat Mumtaz Virk, Lars Borin, Anju Saxena, and	<a href="#">in-context linguistic descriptions</a> . <i>arXiv preprint</i>	974
923	Harald Hammarström. 2017. <a href="#">Automatic extraction</a>	arXiv:2402.18025.	975
924	<a href="#">of typological linguistic features from descriptive</a>		
925	<a href="#">grammars</a> . In <i>Text, Speech, and Dialogue: 20th In-</i>		
926	<i>ternational Conference, TSD 2017, Prague, Czech</i>		
927	<i>Republic, August 27-31, 2017, Proceedings 20</i> , pages		
928	111–119. Springer.		
929	Shafqat Mumtaz Virk, Daniel Foster, Azam Sheikh		
930	Muhammad, and Raheela Saleem. 2021. <a href="#">A deep</a>		
931	<a href="#">learning system for automatic extraction of typologi-</a>		
932	<a href="#">cal linguistic information from descriptive grammars</a> .		



## A Rerankers: HuggingFace

1. bge-en-icl (Xiao et al., 2023): <https://huggingface.co/BAAI/bge-en-icl>  
License: Apache 2.0
2. stella\_en\_1.5B\_v5: [https://huggingface.co/dunzhang/stella\\_en\\_1.5B\\_v5](https://huggingface.co/dunzhang/stella_en_1.5B_v5)  
License: MIT
3. NV-Retriever-v1 (de Souza et al., 2024): <https://huggingface.co/nvidia/NV-Retriever-v1>  
License: NVIDIA license agreement
4. gte-Qwen2-7B-instruct (Li et al., 2023): <https://huggingface.co/Alibaba-NLP/gte-Qwen2-7B-instruct>  
License: Apache 2.0
5. Linq-Embed-Mistral (Kim et al., 2024): <https://huggingface.co/Linq-AI-Research/Linq-Embed-Mistral>  
License: CC-BY-NC-4.0
6. SFR-Embedding-2\_R (Meng\* et al., 2024): [https://huggingface.co/Salesforce/SFR-Embedding-2\\_R](https://huggingface.co/Salesforce/SFR-Embedding-2_R)  
License: CC-BY-NC-4.0
7. SFR-Embedding-Mistral (Meng et al., 2024): <https://huggingface.co/Salesforce/SFR-Embedding-Mistral>  
License: CC-BY-NC-4.0

## B Benchmark for Rerankers: Examples for Relevance Categories

0 — the paragraph does not mention word order at all.

*In order to express ‘from’, these demonstrative members must take the ablative-1 suffix (-ngomay), like all other adverb and also nouns, e.g. yarro-ngomay ‘here-ABL1’ in (3-134), (4-13), (4-18-b), (4-77), (4-114). (Tsunoda, 2012, p. 181)*

Annotated with 0: the paragraph only describes demonstratives without any mentions of orders of elements in the language.

Paragraphs that contain examples with glosses without explicit mentions of word order were similarly annotated with 0: examples without context should not be treated as evidence of a language

having a particular word order, because a language may have no dominant word order.

1 — the paragraph mentions or describes word order in a construction other than the monotransitive construction (or order of morphemes/phonemes/clitics/etc.), since WALS 81A refers to the word order in the monotransitive construction with the verb in the declarative mood in particular.

*Table 36 shows that propositional enclitics are ordered in relation to each other. The directional enclitics -(e)nhdhi ‘TOWARDS’ and -(e)ya ‘AWAY’ are mutually exclusive. (Ford, 1998, p. 269)*

Annotated with 1: the paragraph describes order of enclitics instead of the constituent order in the monotransitive construction.

2 — the paragraph mentions or describes the word order in the monotransitive construction (in a title of a section, in the table of contents, or in references).

*736 27.3 Word order at the clause level . . . . . (Forker, 2013, p. xxiv)*

Annotated with 2: this chunk is a fragment of a table of contents.

3 — the paragraph mentions or describes the word order in the monotransitive construction (in a paragraph in the main text).

*Alsea, Siuslaw, and Coos have been tentatively categorized as having VOS as their basic word order, by Greenberg (1966), on the basis of the fact that Greenberg found VOS to be the most common order of subject, object, and verb in these languages. (Morgan, 1991, p. 482)*

Annotated with 3: the grammar describes the language Kutenai, but this paragraph mentions the constituent order in the monotransitive construction in other languages.

4 — the paragraph narrows down the word order in the monotransitive construction to several variants.

*As a consequence of its predominant verb-medial order, Qaqet does not have any clause chaining and/or switch reference <...> (Hellwig, 2019, p. 19)*



1067 Annotated with 4: the mention that Qaqet has a  
1068 predominantly verb-medial order narrows the seven  
1069 logically possible variants to “SVO”, “OVS”, and  
1070 “No dominant order”.

1071 5 — a linguist can unambiguously determine the  
1072 constituent order in the monotransitive construction  
1073 from the paragraph.

1074 <...> *The constituent order in relative*  
1075 *clauses is SOV, as in main clauses. The*  
1076 *subject in relative clauses is obligatorily*  
1077 *encoded as genitive, while all other con-*  
1078 *stituents appear as they would in an inde-*  
1079 *pendent verbal clause. (Wegener, 2012,*  
1080 *p. 254)*

1081 Annotated with 5: the word order is explicitly  
1082 mentioned in the paragraph.

1083 *One might expect that the peculiar con-*  
1084 *stituent order of Urarina would also be*  
1085 *subject to pressure from Spanish (a noto-*  
1086 *rious A V O / S V language), but signifi-*  
1087 *cant changes to constituent order in Ura-*  
1088 *rina are not observed. As mentioned in*  
1089 *§18.3, there are a few isolated examples*  
1090 *of an S or A argument occurring in pre-*  
1091 *verbal position that cannot be accounted*  
1092 *for in terms of the predicted features (fo-*  
1093 *cus, emphasis, negation). Beside that,*  
1094 *in one of the dialects investigated fur-*  
1095 *ther above (Copal), two examples with*  
1096 *an O argument in postverbal position*  
1097 *were observed. While such examples are*  
1098 *extremely rare, one could of course at-*  
1099 *tribute these to the influence of Spanish.*  
1100 *(Olawsky, 2006, p. 899)*

1101 Annotated with 5. Although there is no explicit  
1102 mention of the word order in Urarina, it is described  
1103 in the paragraph that there are only isolated exam-  
1104 ples of the subject argument in Urarina occurring  
1105 in the preverbal position and of the object argument  
1106 occurring in the postverbal position. Consequently,  
1107 the only possible logical variant that is possible  
1108 for Urarina is OVS, contrary to the immediately  
1109 obvious mention of SVO (A V O) in Spanish. Ex-  
1110 tracting information from such paragraphs based  
1111 solely on term frequency would be suboptimal.

## C Benchmark for Rerankers: Details

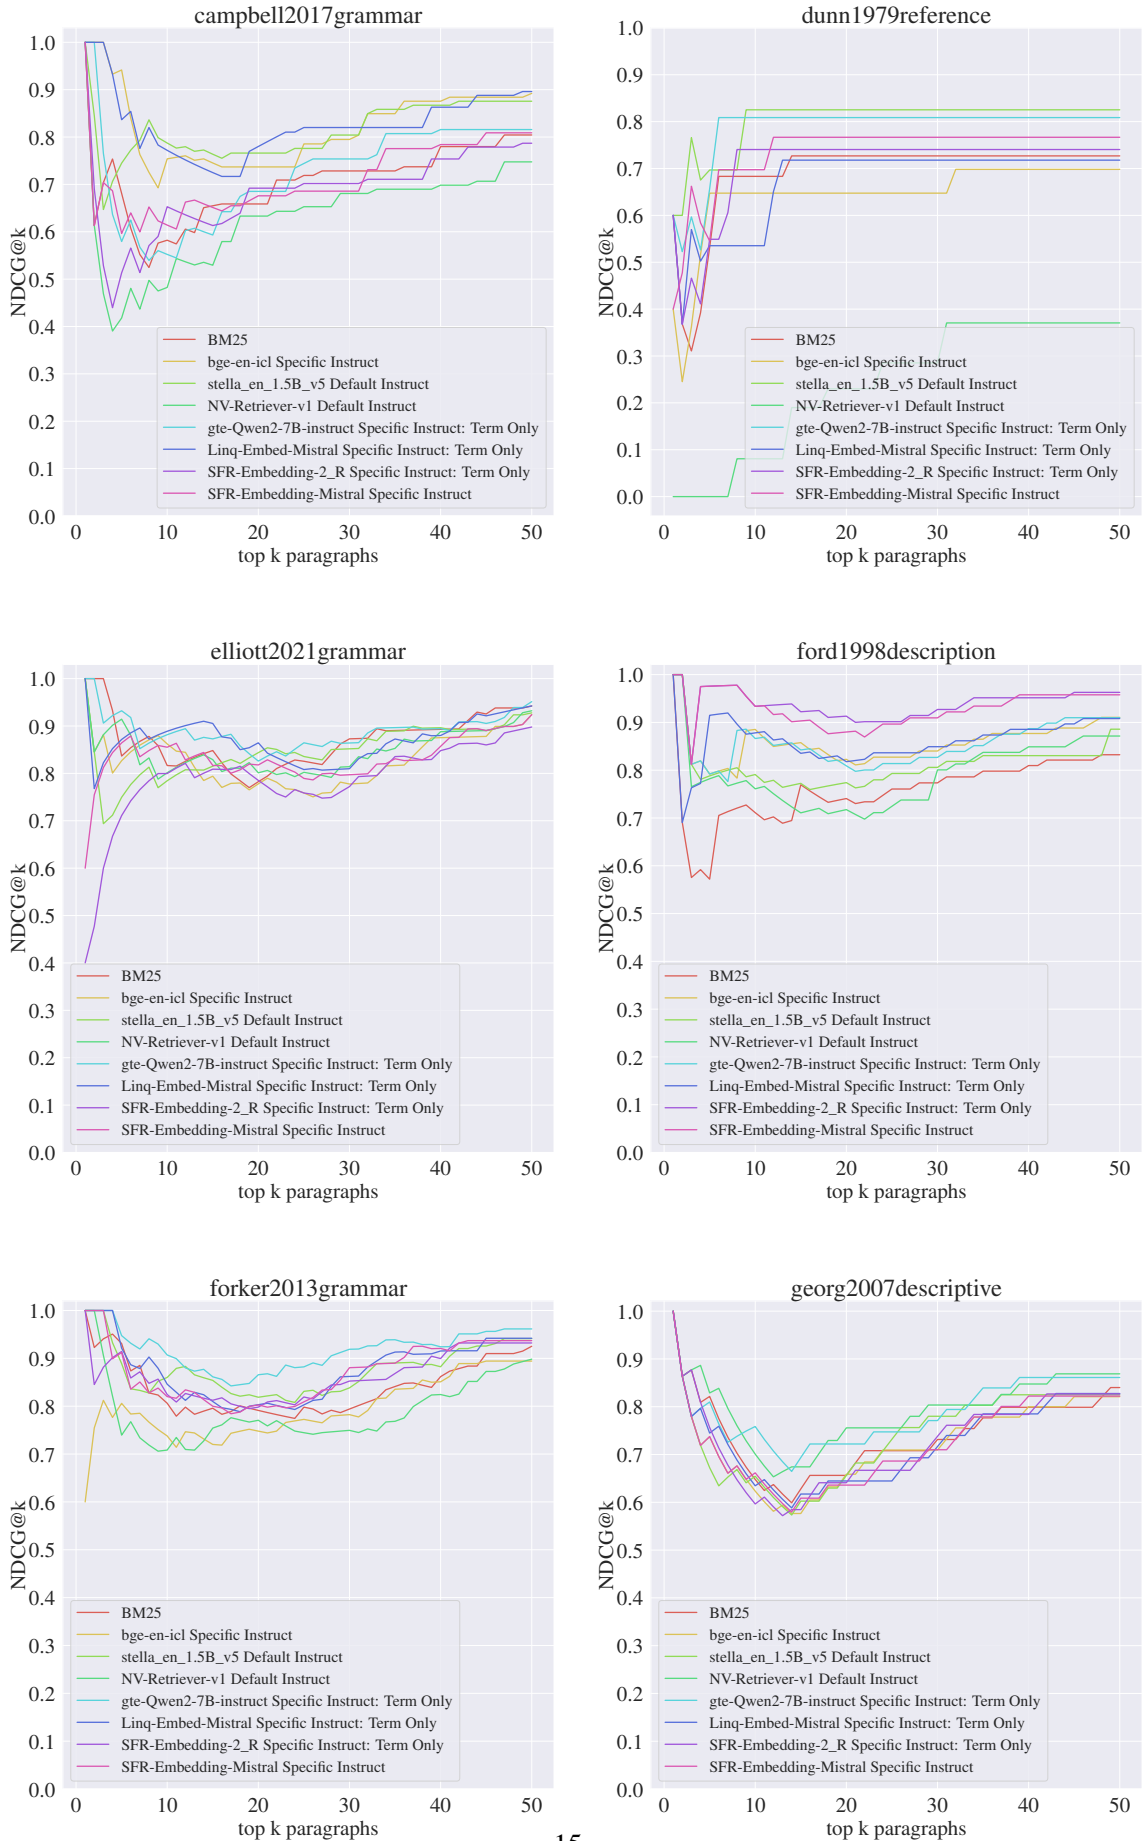
Grammar	Language	Family	Macroarea	Word order
(Campbell, 2017)	Gã	Niger-Congo	Africa	SVO
(Newman, 2002)	Hausa	Afro-Asiatic	Africa	No mention
(Georg, 2007)	Ket	Yeniseian	Eurasia	No mention
(Forker, 2013)	Hinuq	Nakh-Daghestanian	Eurasia	SOV
(Ford, 1998)	Emmi	Western Daly	Australia	No dom. order
(Tsunoda, 2012)	Warrongo	Pama-Nyungan	Australia	SOV
(Morgan, 1991)	Kutenai	Kutenai	North America	VOS
(Dunn, 1979)	Tsimshian (Coast)	Tsimshianic	North America	VSO
(Elliott, 2021)	Enxet Sur	Lengua	South America	No dom. order
(Sakel, 2004)	Mosetén	Mosetenan	South America	SVO
(Hellwig, 2019)	Qaqet	Baining	Papunesia	SVO
(Wegener, 2012)	Savosavo	Solomons East Papuan	Papunesia	SOV
(Olawsky, 2006)	Urarina	Urarina	South America	OVS
(Weir, 1986)	Nadëb	Nadahup	South America	OSV







Table 5: Languages featured in the benchmark for rerankers. “No mention” in the column “Word order” indicates that the 50 paragraphs chosen by BM25 for this grammar do not allow to narrow down the word order in the language to one feature.

Grammar	Annotations						0	1-3	4-5
	0	1	2	3	4	5	(irrelevant)	(misleading)	(relevant)
(Campbell, 2017)	34	7	0	2	0	7	0.68	0.18	0.14
(Newman, 2002)	15	30	3	1	1	0	0.30	0.68	0.02
(Georg, 2007)	36	13	0	0	1	0	0.72	0.26	0.02
(Forker, 2013)	9	14	10	6	4	7	0.18	0.60	0.22
(Ford, 1998)	28	19	0	0	0	3	0.56	0.38	0.06
(Tsunoda, 2012)	8	23	7	5	3	4	0.16	0.70	0.14
(Morgan, 1991)	0	12	10	11	4	13	0.00	0.66	0.34
(Dunn, 1979)	45	0	1	2	0	2	0.90	0.06	0.04
(Elliott, 2021)	5	18	2	4	3	18	0.10	0.48	0.42
(Sakel, 2004)	11	19	5	6	4	5	0.22	0.60	0.18
(Hellwig, 2019)	24	10	1	0	7	8	0.48	0.22	0.30
(Wegener, 2012)	17	16	7	3	1	6	0.34	0.52	0.14
(Olawsky, 2006)	6	17	3	2	8	14	0.12	0.44	0.44
(Weir, 1986)	34	4	2	2	5	3	0.68	0.16	0.16
<b>Total</b>	<b>272</b>	<b>202</b>	<b>51</b>	<b>44</b>	<b>41</b>	<b>90</b>	<b>0.3886</b>	<b>0.4243</b>	<b>0.1871</b>

Table 6: Distribution of annotation categories in the benchmark for rerankers.

## D Benchmark for Rerankers: Dynamics for NDCG@k
















Figure 1: NDCG@k across all values of k for best performing configurations for each model.







One of the most evident phenomena indicated by the plots is low NDCG@k scores at small values of k on (Newman, 2002), indicating that all models ranked an irrelevant paragraph first. The “irrelevant” paragraph in question appears to be the following:

*WORD ORDER The basic word order in sentences with i.o.’s is V + i.o. + (d.o.), i.e., verb followed immediately by the sh. indirect object followed by the direct object (if present). This word order is the same whether the i.o. is a noun or a pronoun, e.g., ura v3amigrate Misa ya kawo [wa tsohuwa];9, ruwaMusa brought the old woman water. (Newman, 2002)*

This paragraph describes the basic word order in a ditransitive construction and has been annotated with 1 (“the paragraph mentions or describes word order in a construction other than the monotransitive construction”). While a linguist may make a rational assumption that the monotransitive construction also follows the VO pattern, it would be incorrect to make assumptions in most similar cases: for instance, WALS lists 13 languages that have the SVO dominant order for transitive constructions, but VS for intransitive constructions, and any newly discovered language may potentially violate a principle previously considered a universal.

## E Benchmark for Rerankers: Wikipedia Summary vs. Term Only








Figure 2: Mean NDCG@k for all grammars: Wikipedia Summary vs Term Only.



## F Benchmark for RAG: Prompts

Prompt for Word Order (partially based on WALS 81A):

*Please determine the dominant word order (order of subject, object, and verb) in the language <...>.*

*The term "dominant word order" in the context of this feature refers to the dominant order of constituents in declarative sentences, in the case where both the subject and the object participants are nouns.*

*Reply with one of the 7 following options: SOV, SVO, VOS, VSO, OVS, OSV, No dominant order.*

*1. Provide the reasoning for the chosen option.*

*2. After the reasoning, output the word "Conclusion:" and the chosen option at the end of your response.*

Prompt for Standard Negation (partially based on GB107):

*Please determine if standard negation in the language <...> can be marked by a modification of the verb or an affix/clitic that is phonologically bound to the verb.*

*The term "standard negation" refers to constructions that mark negation in declarative sentences involving dynamic (not-stative) verbal predicates.*

*Morphemes that attach (become phonologically bound) to other constituents, not verbs only, do not count.*

*Clitic boundaries are marked in the glosses by an equals sign: "=".*

*Affix boundaries are marked in the glosses by a dash: "-".*

*Separate words (i.e., particles that are not phonologically bound to other words) are separated from other words by spaces.*

*Choose one of the 2 following options: 1, 0.*

*Reply with 1 if standard negation in the language <...> can be marked by an affix, clitic or modification of the verb.*

*Reply with 0 if standard negation in <...> cannot be marked by an affix, clitic or modification of the verb.*

*1. Provide the reasoning for the chosen option.*

*2. After the reasoning, output the word "Conclusion:" and the chosen option at the end of your response.*

Prompt for Polar Questions (partially based on Grambank chapters related to strategies for marking polar questions):

*Please determine all possible strategies for forming polar questions (yes-no questions) in the language <...>.*

*Consider neutral polar questions only (non-neutral, or leading, polar questions indicate that the speaker expects a particular response).*

*The 7 strategies for forming polar questions are the following: Interrogative intonation only, Interrogative word order, Clause-initial question particle, Clause-final question particle, Clause-medial question particle, Interrogative verb morphology, Tone.*

*Clitic boundaries are marked in the glosses by an equals sign: "=".*

*Affix boundaries are marked in the glosses by a dash: "-".*

*Separate words (e. g. particles that are not phonologically bound to other words) are separated from other words by spaces.*

*For this feature, count interrogative clitics as particles if they can be bound to other constituents in the sentence, not to the verb only.*

*Interrogative morphemes that can be phonologically bound to the verb only are counted as interrogative verbal morphology.*

*If a morpheme (for example, clitic or particle) can follow any constituent, which can be in various positions within the clause, including the clause-final position, code 1 for both "Clause-medial question particle" and "Clause-final question particle".*

*If a morpheme (for example, clitic or particle) can precede any constituent, which can be in various positions within the clause, including the clause-initial position, code 1 for both "Clause-initial question particle" and "Clause-medial question particle".*

*For each strategy, code 1 if it is present in the described language; code 0 if it is absent in the language.*

*Example of the output for a language that marks polar questions either with interrogative intonation only or with a clause-final interrogative particle:*

*"Interrogative intonation only: 1, Interrogative word order: 0, Clause-initial question particle: 0, Clause-final question particle: 1, Clause-medial question particle: 0, Interrogative verb morphology: 0, Tone: 0"*

*1. Provide the reasoning for the chosen option.*

*2. After the reasoning, output the word "Conclusion:" and the chosen option at the end of your response.*

Prompt for Number of Cases (partially based on WALS 49A):

*Please determine the number of cases in the language <...>.*

*The term "cases" in the context of this feature refers to productive case paradigms of nouns.*

*Reply with one of the 9 following options: No morphological case-marking, 2 cases, 3 cases, 4 cases, 5 cases, 6-7 cases, 8-9 cases, 10 or more cases, Exclusively borderline case-marking.*

*The feature value "Exclusively borderline case-marking" refers to languages which have overt marking only for concrete (or "peripheral", or "semantic") case relations, such as locatives or instrumentals.*

*Categories with pragmatic (non-syntactic) functions, such as vocatives or topic markers, are not counted as case even if they are morphologically integrated into case paradigms.*

*Genitives are counted as long as they do not encode categories of the possessum like number or gender as well, if they do not show explicit adjective-like properties. Genitives that may take additional case affixes agreeing with the head noun case ("double case") are not regarded as adjectival.*

*1. Provide the reasoning for the chosen option.*

*2. After the reasoning, output the word "Conclusion:" and the chosen option at the end of your response.*

## G Benchmark for RAG: Prompt Details

### G.1 Wikipedia Summaries

Feature	Wikipedia Title
WALS 81A	Word order <sup>3</sup>
GB 107	Affirmation and negation <sup>4</sup>
WALS 49A	Grammatical case <sup>5</sup>
WALS 116A*	Yes-no question <sup>6</sup>

Table 7: The Wikipedia pages we chose to take their summaries as definitions for the typological features (versions as of September 3rd, 2024).

The fourth paragraph from the "Grammatical case" summary was not included in the prompts due to containing examples of languages and their numbers of cases. If it was included in the prompt, we would have essentially provided the RAG system with an answer to a part of the given task.

### G.2 Base Prompts

Due to the multilabel classification on the linguistic domain already posing a potential significant challenge for the RAG pipeline, we simplified the annotation guidelines for the features pertaining to interrogative particles compared to the ones presented in Grambank.

Firstly, in our version of the feature, the particles do not necessarily have to be dedicated to marking polar questions; the main purpose of the particle may be marking content questions, as long as using it to mark polar questions is possible as well. The motivation for this adjustment is the fact that the retriever and reranker in the RAG pipeline will retrieve information related to polar questions only based on the query, and it is irrational to expect information on polar questions to appear in the retrieved paragraphs.

Secondly, we deemed it unnecessary for a clause-medial particle to have the middle of the clause as its most common placement; a particle is considered clause-medial if it can appear in the clause-medial position in at least one example.

Another simplification is omission of two other features related to polar questions in Grambank:

<sup>3</sup>[https://en.wikipedia.org/w/index.php?title=Word\\_order&oldid=1240489972](https://en.wikipedia.org/w/index.php?title=Word_order&oldid=1240489972)

<sup>4</sup>[https://en.wikipedia.org/w/index.php?title=Affirmation\\_and\\_negation&oldid=1202783032](https://en.wikipedia.org/w/index.php?title=Affirmation_and_negation&oldid=1202783032)

<sup>5</sup>[https://en.wikipedia.org/w/index.php?title=Grammatical\\_case&oldid=1238822129](https://en.wikipedia.org/w/index.php?title=Grammatical_case&oldid=1238822129)

<sup>6</sup>[https://en.wikipedia.org/w/index.php?title=Yes%280%93no\\_question&oldid=1236424936](https://en.wikipedia.org/w/index.php?title=Yes%280%93no_question&oldid=1236424936)

GB286, GB297. V-not-V constructions and double-marking with both a particle and verbal morphology occur in our benchmark in isolated examples only.

Despite the simplifications, we impose the requirement on assigning clitics to the particle category if they can attach to anything in their vicinity, and to the verb morphology category if they can be phonologically bound to the verb only, in order for the task to test the capabilities of the RAG pipeline in regard to distinguishing morphology from syntax. GB291 (marking polar questions by means of tone) has been retained for a similar purpose: distinguishing morphology from prosody.

### G.3 Modifications for Chain-of-Thought

We modified each chapter of the WALS and Grambank for the Chain-of-Thought prompts in the following way according to our specifications: We eliminated all mentions and examples of languages that already exist in the respective benchmarks. For instance, (1) has been modified to obtain (2) due to Oneida being present in the Number of Cases benchmark:

(1) *This excludes, for example, the "locative" suffixes in Oneida (Iroquoian; Ontario) from being counted as case, since they can derive body-part nouns which may occur in all semantically permitted syntactic positions (i.e. not only as locational adverbials). (WALS, Chapter 49A: Number of Cases)*

(2) *For example, "locative" suffixes that can derive body-part nouns which may occur in all semantically permitted syntactic positions (i.e. not only as locational adverbials), are not counted as case.*

Additionally, we removed mentions of the WALS map, since each prompt for the RAG pipeline is formulated as determining a feature value instead of placing it on a map. We eliminated duplicated information (i.e., left only one mention about disregarding leading polar questions across all chapters in GB257 – GB 291); changed the guidelines according to our simplified definitions of clause-initial/clause-medial/clause-final particles; and included the explanation of when clitics are regarded as particles and when they are considered verbal morphology.



## H Benchmark for RAG: Distribution of Values

Word Order (WALS 81A)	
No dominant order	40
No mention	12
OSV	1
OVS	1
SOV	59
SVO	25
VOS	4
VSO	6
<b>Total</b>	148

Number of Cases (WALS 49A)	
No case marking	55
Excl. borderline case marking	21
2 cases	11
3 cases	3
4 cases	5
5 cases	6
6-7 cases	22
8-9 cases	13
10 or more cases	12
<b>Total</b>	148

Standard Negation (GB 107)	
0	65
1	83
<b>Total</b>	148

Polar Questions (WALS 116A*)	
Clause-initial particle	
0	20
1	128
Clause-medial particle	
0	132
1	16
Clause-final particle	
0	100
1	48
Verb morphology	
0	125
1	23
Interrogative word order	
0	144
1	4
Interrogative intonation only	
0	96
1	52
Tone	
0	141
1	7
<b>Total</b>	148

Table 8: Distribution of values in each feature (total = 148). The imbalance of classes should not be mitigated, because the benchmark, serving as a representative language sample, is intended to reflect the natural distribution of values found in languages around the world. If the grammar does not mention the feature, the following values are assigned to it: *No mention* for **WALS 81A**; *No case marking* for **WALS 49A**; *0* for **GB 107**; *0* for each of the sub-features for **WALS 116A\***. We make the assumption that a lack of mention in the descriptive material most likely indicates a lack of the phenomenon in the language (apart from the Word Order feature, since one of the seven logically possible categories must inevitably be present).

## I Benchmark for Rerankers: Ablation

Additionally to the main experiment, we conducted an ablation study: how would the RAG pipeline perform if it had access to a near-optimal retriever instead of the combination of BM25 and rerankers, i.e., how much of the error is attributable to the generation component of Retrieval-Augmented Generation? In order to conduct this ablation study, we annotated the pages in the grammar documents where the relevant information for the feature is contained (i.e., sufficient information for a linguist to unambiguously infer the value of the feature). Since this approach is only applicable to grammars that actually contain the relevant information, we ran the ablation on 136 grammars for WALS 81A, 146 grammars for GB 107, 121 grammars for WALS 116A, and 140 grammars for WALS 49A (out of 148), the same subsets as for the baseline runs.

We used the manually retrieved pages as the retrieved component in RAG, resulting in two columns on the right in Table 9, and recalculated the metrics for the four RAG configurations on the same subsets as used for the ablation. The baseline metrics are identical to the ones in Table 4 due to identical subsets of grammars.

The results indicate that for features that have a high baseline performance (most likely due to having an extreme class imbalance), e.g., *Interrogative Word Order* and *Tone*, involving a human to retrieve the paragraphs may harm the classification quality.

For WALS 49A (the feature hypothesized to be the most difficult one to classify) it is the Chain-of-Thought component that makes the difference between harming the classification performance and improving it; however, the F1 scores for this feature remain low, not exceeding 0.56.

	F1 average	Baseline	Wiki	Wiki +CoT	Re- ranker	Rer. +CoT	Human	Human +CoT
WALS 81A	micro	0.5551 ± 0.0359	0.7353	0.7426	0.7500	<b>0.7794</b>	<b>0.7868</b>	<b>0.7868</b>
	macro	0.2812 ± 0.0276	<b>0.7060</b>	0.5866	0.5880	0.6000	<b>0.5198</b>	0.4949
	weighted	0.5328 ± 0.0337	0.7398	0.7499	0.7391	<b>0.7721</b>	0.7838	<b>0.7939</b>
GB 107	micro	0.5747 ± 0.0350	0.6644	0.6644	0.6849	<b>0.7192</b>	0.8151	<b>0.8425</b>
	macro	0.5740 ± 0.0354	0.5679	0.5814	0.6100	<b>0.6644</b>	0.8038	<b>0.8304</b>
	weighted	0.5724 ± 0.0361	0.5958	0.6069	0.6334	<b>0.6829</b>	0.8102	<b>0.8366</b>
WALS 49A	micro	0.3986 ± 0.0238	0.5000	0.5214	<b>0.5357</b>	0.5000	0.5000	<b>0.5429</b>
	macro	0.2216 ± 0.0158	<b>0.4655</b>	0.4461	0.4299	0.4031	0.4126	<b>0.4737</b>
	weighted	0.3453 ± 0.0237	0.5088	0.5341	<b>0.5409</b>	0.5023	0.5082	<b>0.5579</b>
interrog. intonation only	micro	0.4504 ± 0.0396	0.8347	<b>0.8926</b>	0.8430	0.8678	0.9174	<b>0.9256</b>
	macro	0.4240 ± 0.0456	0.8346	<b>0.8924</b>	0.8429	0.8677	0.9169	<b>0.9251</b>
	weighted	0.4068 ± 0.0484	0.8352	<b>0.8929</b>	0.8433	0.8681	0.9178	<b>0.9260</b>
interrog. word order	micro	0.9653 ± 0.0076	<b>0.9917</b>	0.9835	0.9835	0.9752	<b>0.9917</b>	0.9835
	macro	0.6641 ± 0.0460	<b>0.9423</b>	0.8957	0.8707	0.8269	<b>0.9264</b>	0.8707
	weighted	0.9611 ± 0.0064	<b>0.9922</b>	0.9851	0.9835	0.9765	<b>0.9912</b>	0.9835
clause- initial particle	micro	0.7264 ± 0.0402	0.9008	0.9091	<b>0.9421</b>	0.9256	<b>0.9587</b>	0.9339
	macro	0.5446 ± 0.0502	0.8127	0.8385	<b>0.8972</b>	0.8679	<b>0.9266</b>	0.8889
	weighted	0.7371 ± 0.0337	0.8987	0.9100	<b>0.9427</b>	0.9263	<b>0.9591</b>	0.9362
clause- final particle	micro	0.4909 ± 0.0409	0.7190	0.7851	0.7025	<b>0.7851</b>	<b>0.8017</b>	0.8017
	macro	0.4808 ± 0.0415	0.7190	0.7826	0.7025	<b>0.7826</b>	<b>0.8000</b>	0.7977
	weighted	0.4661 ± 0.0428	0.7195	0.7874	0.7030	<b>0.7874</b>	<b>0.8038</b>	0.8035
clause- medial particle	micro	0.6760 ± 0.0367	0.8430	<b>0.8843</b>	0.8430	0.8595	<b>0.9091</b>	0.9008
	macro	0.5011 ± 0.0544	0.7961	<b>0.8374</b>	0.7961	0.8049	<b>0.8671</b>	0.8606
	weighted	0.6644 ± 0.0355	0.8502	<b>0.8857</b>	0.8502	0.8621	<b>0.9085</b>	0.9021
interrog. verb morphology	micro	0.7678 ± 0.0185	0.7686	<b>0.8099</b>	0.8017	<b>0.8182</b>	0.8760	<b>0.8843</b>
	macro	0.4533 ± 0.0311	0.7026	0.7361	0.7279	<b>0.7506</b>	0.8188	<b>0.8235</b>
	weighted	0.7102 ± 0.0160	0.7894	0.8226	0.8157	<b>0.8311</b>	0.8819	<b>0.8877</b>
tone	micro	0.9273 ± 0.0128	0.9174	0.9504	0.9587	<b>0.9669</b>	0.9504	<b>0.9669</b>
	macro	0.5147 ± 0.0573	0.7500	0.8199	<b>0.8572</b>	0.8484	0.8199	<b>0.8662</b>
	weighted	0.9104 ± 0.0121	0.9309	0.9555	0.9637	<b>0.9669</b>	0.9555	<b>0.9689</b>

Table 9: F1 scores for the ablation experiments (on the right) and the four initial RAG configurations + baseline (on the left) on the same subset of grammars as the ablation experiments. The best F1 value in each row in the main table and the best F1 value for the better of the two ablation experiments are shown in **bold**. Each of the two ablation experiments was executed only once. Blue cells in the ablation experiments indicate **better** performance, red cells indicate **worse** performance, and purple cells indicate **identical** performance compared to the the best non-human retriever in the same row.

Most features demonstrate an increase in F1 scores when compared to non-human retrievers.

The greatest improvement in quality was achieved on the balanced binary value, GB 107: +0.166 macro F1 compared to the best RAG configuration.

Overall, a conclusion can be made that while the task of machine reading on the linguistic domain is still not fully resolved, the task of information retrieval within the linguistic domain is similarly important.