

# A Rigorous Evaluation of LLM Data Generation Strategies for Low-Resource Languages

Tatiana Anikina<sup>♣\*</sup>, Jan Cegin<sup>†‡\*</sup>, Jakub Simko<sup>‡</sup>, Simon Ostermann<sup>♣</sup>

♣ German Research Institute for Artificial Intelligence (DFKI), Saarbrücken, Germany

{tatiana.anikina, simon.ostermann}@dfki.de

† Faculty of Information Technology, Brno University of Technology, Brno, Czechia

‡ Kempelen Institute of Intelligent Technologies, Bratislava, Slovakia

{jan.cegin, jakub.simko}@kinit.sk

## Abstract

Large Language Models (LLMs) are increasingly used to generate synthetic textual data for training smaller specialized models. However, a comparison of various generation strategies for low-resource language settings is lacking. While various prompting strategies have been proposed—such as demonstrations, label-based summaries, and self-revision—their comparative effectiveness remains unclear, especially for low-resource languages. In this paper, we systematically evaluate the performance of these generation strategies and their combinations across 11 typologically diverse languages, including several extremely low-resource ones. Using three NLP tasks and four open-source LLMs, we assess downstream model performance on generated versus gold-standard data. Our results show that strategic combinations of generation methods—particularly target-language demonstrations with LLM-based revisions—yield strong performance, **narrowing the gap with real data to as little as 5% in some settings**. We also find that smart prompting techniques can **reduce the advantage of larger LLMs**, highlighting efficient generation strategies for synthetic data generation in low-resource scenarios with smaller models.

## 1 Introduction

Large language models (LLMs) such as GPT-4, Gemini, and Llama show impressive performance in producing well-formed text - trivially. This essential capability makes them ideal for both the *augmentation* and *generation* of text datasets (Ubani et al., 2023; Dai et al., 2023; Piedboeuf and Langlais, 2023; Cegin et al., 2023, 2025). Such LLM-based text generation is often leveraged for creating representative data that are used for training smaller, more efficient downstream models, a technique that is sometimes referred to as *LLM distillation* (Xu et al., 2024). LLM-based

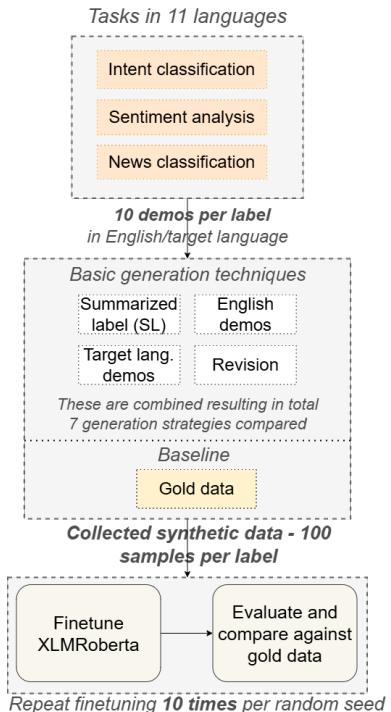


Figure 1: Overview of our methodology. We consider 11 different languages (including English) for 3 different tasks. We use various generation strategies to create synthetic data for each strategy, language and LLM combination. XLM-R is then finetuned and evaluated on the test set of the original dataset for the given task and language.

text generation has been used for various downstream tasks such as sentiment analysis (Onan, 2023; Piedboeuf and Langlais, 2023), intent recognition (Cegin et al., 2024b) and news topic classification (Piedboeuf and Langlais, 2023; Cegin et al., 2024a). While most studies focus only on English, LLMs have also been leveraged for textual generation in low-resource languages in some cases (Namboori et al., 2023; Glenn et al., 2023; Zotova et al., 2021; Chung et al., 2025), showing impressive performance in generating textual data for a wide variety of languages and often outperforming neural machine translation as augmentation.

\*Equal contribution.

tion technique (Pranida et al., 2025).

Various generation strategies exist for generating high-quality textual data, such as including textual examples of the desired data in the generation prompt itself (Cegin et al., 2024a), using descriptions of the labels of data to be generated in the generation prompt (Gu et al., 2024; Abraham et al., 2025) and using self-revision of the LLMs after generation to filter generated data (Li et al., 2024). **However, two essential gaps remain.** First, the performance of such generation strategies is often evaluated in isolation, and their combinations have not been compared, leaving the whole potential of LLM-based data augmentation obscure. Second, they have been used for generating mostly English texts, which leaves most languages under-explored - especially low-resource ones. To the best of our knowledge, an overview of which generation strategies (or their combination) are best for textual generation in low-resource languages is completely lacking.

The goal of this paper is to close this gap and compare existing generation strategies for low-resource text generation in various languages. Our methodology is shown in Figure 1. We evaluate the performance of strategies based on the downstream task performance, using smaller models finetuned on the generated data. We perform a comparative analysis of different generation strategies using 11 typologically diverse languages (including English) with different scripts, including several very low-resource languages such as Welsh, Romanian, Azerbaijani, Slovenian, and Telugu. As generation strategies, we consider 3 prominent approaches: *Summarized label* - using explanations of the data to be generated (Gu et al., 2024; Abraham et al., 2025); *Demonstrations* - the inclusion of examples in the target language or English in the prompts (Piedboeuf and Langlais, 2023; Cegin et al., 2024a); And *Revision* - additionally filtering out samples that are generated using either of the previous two strategies, using the same LLM (Li et al., 2024). We benchmark four open-source LLMs of different sizes on three NLP tasks and perform an ablation study with respect to the language of demonstrations, intent description, and LLM-based revision. We measure the performance of downstream models for various generation strategies and their combinations. For each task, we compare the configurations to an upper bound that is set by the same model being trained on the same amount of real data. All generated data

will be released upon acceptance and our code is available at <https://github.com/tanikina/multilingual-generation>.

Our main contributions are:

- We provide an exhaustive evaluation of different common strategies for **generating synthetic data for low-resource languages** and formulate suggestions for the most effective combination of strategies for extreme low-resource settings.
- We confirm that while models with a higher number of parameters outperform their smaller counterparts on the generation task in the low-resource setting, **the gap in performance is small** when a right generation technique is used.
- We show that - using the right combination of techniques, and for some configurations - **the drop of performance for a model trained on LLM-generated data is as small as up to only 5%** absolute when compared to finetuning on the same amount of “real” data.
- We show that a **combination of demonstrations in the target language with LLM-based revisions** generally leads to the best performance across most languages, especially in extremely low-resource settings.

## 2 Related Work

Soon after their advent, new LLMs, such as GPT-4 or Llama started to be used as data augmentation or data generation tools. In general, this newly generated data is used to train smaller downstream models, for better efficiency. LLM-based augmentation is typically done through paraphrasing (Cegin et al., 2024a; Dai et al., 2023; Sen et al., 2023; Evuru et al., 2024). Less often, LLMs are used to create semantically new samples adhering to a given label (Ubani et al., 2023; Cegin et al., 2024b). LLM-based augmentation and data generation has been used for a variety of tasks such as automated scoring (Fang et al., 2023), low-resource language generation (Ghosh et al., 2023), intent classification (Sahu et al., 2022), sentiment analysis (Piedboeuf and Langlais, 2023; Onan, 2023; Yoo et al., 2021), hate speech detection (Sen et al., 2023), news classification (Piedboeuf and Langlais, 2023), content recommendation (Liu et al., 2024), and health symptoms classifications (Dai et al., 2023).

While the main focus in previous research is on creating textual data for English, some works have also explored synthetic data generation for other languages. Machine Translation-based techniques leverage multilingual models for enhancing textual data in languages such as Vietnamese (Feng et al., 2021). Variants of this approach use chain-of-thought prompting with LLMs (Son et al., 2025). Multilingual synthetic generation using LLMs has been leveraged for various tasks like QA (Kramchaninova and Defauw, 2022; Namboori et al., 2023), fact-checking (Chung et al., 2025), NER (Liu et al., 2021), sentiment stance detection (Zotova et al., 2021) and classification (Glenn et al., 2023) - while not focusing on low-resource languages particularly. A study comparing human-created data, machine translation-created data, and LLM-generated synthetic data for culturally nuanced commonsense reasoning in low-resource languages has found that using LLM-generated data can be better than machine translation for downstream classification (Pranida et al., 2025).

While the usage of LLM-based data generation is growing for non-English languages, a comparison of how such synthetic data should be generated, especially for low-resource languages, is lacking.

### 3 Methodology and Experiments

Our methodology is visualized in Figure 1 and follows a simple procedure: For a given task, we extract 10 demos per label in a target language. We then use a range of generation strategies to collect 100 synthetic data points per label. Finally, we train an XLM-R model (Conneau et al., 2019) on the data. Our chose XLM-R because it represents an encoder-only, efficient and small model (279M parameters) that supports 94 languages. We also performed pilot experiments with mBERT (Devlin et al., 2019) but found that XLM-R achieves better scores on average, therefore we focused on the XLM-R downstream model for the final evaluation.

In this project we follow a controlled setting with the same amount of samples per label, i.e. 10 labels for the intent recognition task results in 1,000 synthetic samples, 7 labels for topic classification in 700 samples, and 2 labels for sentiment analysis in 200 samples. In order to address potential overfitting, we use early stopping with a patience of 5 epochs and although our training and development data are from the generated pool, all test sets are always from the original gold dataset, which

prevents the model from performing too well by memorizing training samples. To ensure a robust comparison of various generation strategies, we use 11 different languages (including English), 3 tasks and 4 different LLMs of varying sizes.

#### 3.1 Generation Strategies

We evaluate three strategies: *Summarized Label*, *Demonstrations*, and *Revision*, which are outlined below.

**Summarized Label.** For this strategy, we generate label descriptions by prompting Llama3-70b with 10 English samples per label (see Figure 8 for the prompt and Section A.7 for the summarized label examples). The generated intent summary is then added to the final prompt. Although many labels in the datasets have self-descriptive names (e.g. *geography*, *cooking\_recipe*), adding a summarized description may provide additional information that is useful for text generation.

**Demonstrations.** We compare two settings: Using English examples of the target label (from which we want to generate samples) or using target language examples of the target label. This way we can assess the impact of the target language demonstrations and simulate a scenario when resources for non-English languages are unavailable. To simulate a low-resource scenario for all cases, we limit the number of demonstrations to 10. We provide only the demonstrations from the class we want to generate and select them randomly, following (Cegin et al., 2024b).

**Revision.** We use an LLM to revise and filter out “bad examples” by prompting with the intent description and previously generated samples. This strategy is used in combination with the previous two, as it requires generated data to be used. Figure 2 provides an example of samples that were accepted or rejected by the judge LLM. Figure 7 shows the prompt for LLM-based assessment of the generated samples.

**Additional Combinations.** We also consider two additional combinations of the generation strategies: We combine the summarized label descriptions with demonstrations in English and the target language without revision in both cases. As *revision* has the highest cost of all strategies (additional inference of the LLM), its exclusion in these combinations enables us to measure what impact on downstream model performance revision has.

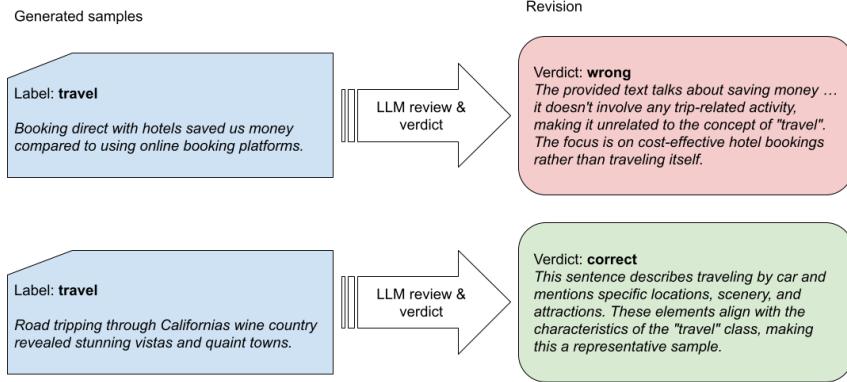


Figure 2: LLM-based revision examples.

### 3.2 Data and Tasks

We consider three different tasks: intent recognition, topic classification, and sentiment analysis. We generate artificial data for each of these tasks in 11 typologically diverse languages using the generation strategies outlined above. We use 2 high-resource languages, German and English, as a comparison. Our experiments include mid-resourced languages: Thai, Hebrew, Indonesian, and Swahili; as well as low-resourced ones: Romanian, Azerbaijani, Slovenian, Telugu, and Welsh. This selection of languages is based on their diversity and the availability of data for each of the tasks.

For intent recognition, we use MASSIVE (FitzGerald et al., 2023), which is a multilingual dataset for virtual assistant evaluation with intent and slot annotations that supports 51 languages and 60 intents. Given that our focus on data generation in low-resource settings introduces additional complexities, we simplify the intent recognition task by considering only ten distinct labels from the original dataset (see A.4 for the examples).

For topic classification, we chose SIB-200 (Adeleani et al., 2024) that has 7 labels (see A.5 for the examples) and also supports a variety of high and low-resource languages. SIB-200 is based on the Flores-200 machine translation corpus, and it uses sentence-level annotations.

As the third task, we selected sentiment classification. To the best of our knowledge, there is currently no unified dataset available for this task that contains a wide range of high and low-resource languages. Thus, in our experiments, we consider 10 different datasets for low-resource languages from (Gurgurov et al., 2025a) and (Gurgurov et al., 2024) and two additional datasets for English and German from (Mollanorozy et al., 2023). Note that

these datasets substantially vary in terms of topic coverage and text format. For example, the German dataset contains mostly sentiment with respect to transportation and infrastructure, while the Romanian dataset is about product reviews. We acknowledge that although all samples in these datasets express positive or negative sentiment, some topics might be an easier target for text generation.

### 3.3 Models

We generate artificial samples with 4 models of different sizes: Gemma-3 (Team, 2025) with 4 and 27 billion parameters, and Llama-3 (AI@Meta, 2024) with 8 and 70 billion parameters. These models were chosen for their open-source nature and support for multiple languages. Although the extent to which each language was represented in the training data is unclear, we were able to generate samples for all 11 target languages.

### 3.4 Experiments

We generate 100 samples per label for each dataset and language using the generation strategies described in 3.1 (see A.3 for the details on computational resources). After collecting the data, we fine-tuned XLM-R (Conneau et al., 2019) on generated data for all three tasks and eleven languages. As an upper bound for comparison, we also train the model on the same amount of balanced gold data, i.e. 100 examples per class. We fine-tune and evaluate 10 models for each configuration and report macro F1 scores averaged across 10 runs. Details regarding the downstream model fine-tuning can be found in A.6.

Intent Recognition	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	91.62	91.46	94.38	95.32	92.17	92.78	94.27	93.66	89.75	90.62	94.27	92.76
Summarized Label (SL)	78.13	63.22	88.06	92.42	78.54	90.00	89.01	88.47	71.85	81.85	86.35	82.53
EnglishDemos + SL	82.25	68.96	92.08	92.79	83.53	91.30	89.22	85.99	71.06	79.66	86.26	83.92
EnglishDemos + Rev.	83.72	67.93	92.40	<b>93.03</b>	84.10	<u>91.73</u>	90.50	87.18	71.87	83.12	86.83	84.77
TargetDemos	79.91	68.45	90.35	90.54	<b>88.92</b>	90.22	90.11	89.38	78.14	83.36	88.33	85.25
TargetDemos + SL	<u>85.70</u>	<u>73.55</u>	91.77	<u>92.79</u>	88.55	91.02	<b>91.58</b>	<b>90.70</b>	79.10	85.11	<u>88.77</u>	87.15
TargetDemos + Rev.	84.50	72.03	91.00	92.35	<u>88.70</u>	91.06	<u>91.18</u>	89.91	<u>79.13</u>	<u>85.65</u>	88.46	86.72
TargetDemos + SL + Rev.	<b>85.88</b>	<b>77.15</b>	<b>92.56</b>	<b>93.03</b>	88.59	<b>91.85</b>	90.63	<u>89.95</u>	<b>80.44</b>	<b>86.07</b>	<b>89.03</b>	<b>87.74</b>
Topic Classification	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	84.85	77.03	84.45	89.99	83.65	87.77	86.53	86.04	76.64	80.29	86.94	84.02
Summarized Label (SL)	75.96	56.10	72.68	72.89	68.70	72.40	73.27	72.69	59.77	67.62	73.92	69.64
EnglishDemos + SL	76.60	60.04	74.69	78.25	71.56	77.29	73.25	75.32	63.23	66.22	<u>76.48</u>	72.08
EnglishDemos + Rev.	<u>77.95</u>	63.70	75.13	<b>79.35</b>	71.82	<u>78.58</u>	<u>74.89</u>	<u>77.53</u>	<b>66.38</b>	68.00	<b>77.58</b>	<u>73.72</u>
TargetDemos	77.82	56.84	72.99	<u>78.58</u>	<u>71.93</u>	78.38	73.91	77.07	61.02	71.41	73.71	72.15
TargetDemos + SL	73.17	57.54	74.68	78.25	71.16	77.53	73.60	77.16	63.27	70.81	75.79	72.09
TargetDemos + Rev.	<b>78.75</b>	<b>66.23</b>	<b>75.64</b>	<b>79.35</b>	71.15	<b>79.77</b>	<b>75.77</b>	<b>77.80</b>	<u>63.30</u>	<b>72.24</b>	74.81	<b>74.07</b>
TargetDemos + SL + Rev.	76.19	<u>64.13</u>	<u>75.32</u>	76.99	<b>72.07</b>	77.81	73.69	76.33	63.02	<u>71.63</u>	76.27	73.04
Sentiment Analysis	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	71.69	58.77	65.84	80.97	82.23	90.39	90.48	83.40	75.07	82.51	80.14	78.32
Summarized Label (SL)	61.66	40.30	65.88	70.23	<b>75.76</b>	75.49	<u>82.83</u>	50.24	46.45	66.59	72.44	64.35
EnglishDemos + SL	<b>67.71</b>	42.46	65.91	<b>77.18</b>	70.10	83.34	82.33	59.52	43.63	64.13	<b>77.31</b>	66.69
EnglishDemos + Rev.	66.33	41.52	<b>67.85</b>	67.58	67.52	80.95	80.96	59.96	50.82	66.05	76.04	65.96
TargetDemos	64.21	<u>48.35</u>	65.09	<u>76.62</u>	70.38	83.99	80.47	66.22	<b>62.02</b>	65.93	73.14	68.77
TargetDemos + SL	<u>67.56</u>	42.53	65.37	70.23	<u>75.47</u>	<b>86.99</b>	79.76	71.64	48.26	<u>73.11</u>	75.63	68.78
TargetDemos + Rev.	67.01	<b>54.41</b>	<u>66.84</u>	<b>77.18</b>	74.71	85.18	<b>83.34</b>	72.47	<u>60.34</u>	72.35	74.10	<b>71.63</b>
TargetDemos + SL + Rev.	67.19	37.15	63.65	65.57	75.46	<u>86.82</u>	79.74	<b>76.44</b>	50.34	<b>74.99</b>	76.89	68.57

Table 1: Fine-tuning results for the gold and artificial data averaged over 4 generative LLMs on the three tasks. For each configuration we fine-tune ten downstream XLM-R models and report average F1 scores. We bold the best generation strategy in each column and underline the second best. See A.2 for language abbreviations.

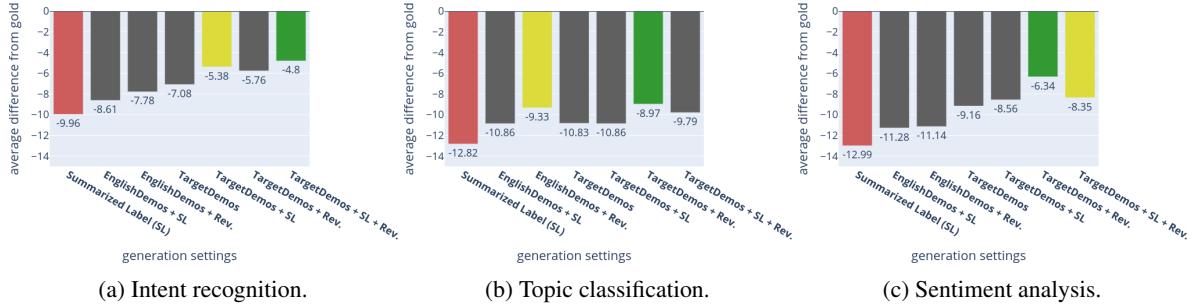


Figure 3: Average difference in performance on the generated data compared to the gold samples across the task dimension. Smaller bars indicate better performance. Green bars show the best-performing generation setting across all languages and models for each task, and yellow bars show the second-best generation setting.

## 4 Results and Discussion

### 4.1 Overview

Table 1 summarizes the results obtained after fine-tuning XLM-R on the LLM-generated data for each of the tasks. The first row per task indicates the scores that can be achieved after fine-tuning on the same amount of gold data, and all subsequent rows correspond to one of the generation strategies. Scores are averaged across four different models,

see Table 9, 10, 11, and 12 in the appendix that show more detailed downstream evaluation results per model, task, and language.

In order to better analyze the quality of samples generated with different strategies for downstream model performance, we consider the three dimensions for a more fine-grained evaluation in the next Sections<sup>1</sup>: Task (4.2), Model (4.3), and Language

<sup>1</sup>For each dimension, we summarize results across other dimensions as a difference to the gold samples rather than

(4.4) dimension. Additionally, in Section 4.5, we analyse in more depth the revision strategy that proved to be most beneficial in our experiments, consistently bringing improvements across all dimensions.

## 4.2 Task Dimension

Figure 3 summarizes the average difference<sup>2</sup> in performance on three tasks for different generation strategies with the blue line on the top representing the upper bound, i.e. the fine-tuning results on the gold data. Ideally, good samples should result in very similar downstream scores. Thus, we expect the best strategy to have the smallest bar (indicated in green, while yellow is the second-best one).

**Intent Recognition.** For the intent recognition task (see Figure 5a) **the best generation strategy includes 10 demonstrations in the target language together with a summarized label description and revision check.** XLM-R fine-tuned on the data generated in this way underperforms the model that was fine-tuned on the gold data only by 4.8% on average. The worst generation setting for the intent classification task is when the LLM is provided only with a summarized label description without any demonstrations. However, if the summarized label description is supplemented with target demonstrations, LLMs can generate good-quality samples that result in an improvement of 4.6% on average. Figure 3 also shows that the revision check helps with both types of demonstrations: For English as well as in the target language.

**Topic Classification.** Using generated samples for topic classification results in a relatively large gap to the gold baseline (up to -12.8% on average, see Figure 3b). However, here we also observe a consistent trend that revision improves generation in all settings. **The best-performing configuration for the topic classification task includes target language demonstrations with revision,** and if demonstrations are replaced with some examples in English, the performance remains very similar ( $\approx 0.4\%$  difference). Notably, summarized label description is detrimental to this task, which might be due to the fact that the class labels are already quite descriptive (e.g. geography, politics), and adding summarized label descriptions may intro-

absolute values, for better readability.

<sup>2</sup>We list the differences as percentage points (absolute difference between F1 scores), not to be confused with percentual (relative) difference.

duce some noise. Besides, English demonstrations are almost as effective as demonstrations in the target language for SIB-200 which can be potentially explained by the fact that SIB-200 data were originally translated from English.

### 4.2.1 Sentiment Analysis

The sentiment analysis task represents the most diverse setting because each language has a separate dataset for evaluation that was not created as part of a unified approach like SIB-200 or MASSIVE. Although all samples are related to sentiment, texts may describe different topics, and the average length of samples differs depending on the language<sup>3</sup>. Nevertheless, this task also demonstrates the advantage of the revision strategy. Moreover, unlike the topic classification task, sentiment analysis substantially benefits from demonstrations in the target language (see Figure 5b).

## 4.3 Model Dimension

Figure 9 illustrates the performance across all tasks and languages for Gemma-3 and Llama-3 models. As expected, **larger models tend to outperform their smaller versions**, e.g. the gap between the gold baseline and the best-performing setting is -6.81% for Llama3-8b and -5.57% for Llama3-70b. Also, Llama models demonstrate slightly better performance across 3 tasks compared to Gemma. Overall, target language demonstrations with revision represent the best generation setting for all models, and adding only summarized label descriptions is not enough to learn useful patterns. Another interesting observation is that Llama benefits more from English demonstrations than Gemma. This could be attributed to the way how multilingual transformers from the Llama family construct latent representations, according to (Wendler et al., 2024), their abstract “concept space” lies closer to English than to other languages, possibly making examples and reasoning in English more useful.

## 4.4 Language Dimension

The main objective of our experiments is to identify the best generation strategies for a low-resource scenario when we have very little (if any) annotated gold data. However, the low-resource status applies not only to the amount of training data available

<sup>3</sup>For instance, gold sentiment examples in Slovenian have longer text (median value is 275 tokens) than Romanian examples (37 tokens).

for the task but also to the target language. Models are expected to be better at generating data in high-resource languages like English or German because they were pre-trained on a large amount of data in those languages. In fact, the results of Figure 10 in the Appendix indicate that English and German achieve relatively good scores in all generation settings, and their best scores are very close to the gold baseline (-3.1% for German and -5.8% for English on average) – synthetic data work almost as good as original data.

However, for lower-resource languages, the performance gap gets larger, e.g., for Welsh, even the best generation setting underperforms the gold baseline by -11.53% on average. Slovenian, Indonesian, Hebrew, Thai, and Telugu achieve the best scores **with a combination of target language demonstrations, intent description and revision**. Azerbaijani, Welsh, Romanian, Swahili and English achieve the highest scores with target language demonstrations and revision but without any intent description. Interestingly, German is the only language that shows slightly better performance when demonstrations are in English. These findings are in line with previous work that attributes an often low performance to LLMs on extremely low-resource languages (Gurgurov et al., 2025b).

English demonstrations without revision perform the worst for Swahili and Telugu, which can most likely be addressed to the lack of pretraining data of such languages. For certain languages (e.g. Indonesian, Thai), XLM-R achieves comparably high F1 scores. This is most likely due to the amount of data in these languages being more prominent in the pretraining corpus of XLM-R (Conneau et al., 2019).

We also analyse the effect of the number of seed samples per label collected on model performance for two of the low-resource languages, Azerbaijani and Welsh in Figure 4. We compare the best-generation strategies with the worst (summarized label) and gold as our baseline. The generation strategies tend to stagnate in performance around 50-60 samples per label collected while adding more gold data generally increases performance.

In summary, **using target language demonstrations together with LLM-based revision yields the best results across both task and model dimensions**. This is also the case for most languages: This combination of generation strategies is best for 4 out of 11 languages and second-best for 5 out of 11 languages. Certain exceptions exist: for

languages like Thai or Hebrew, this combination is not even the second-best generation strategy. Additionally, we note, that even a small amount of demonstrations given to the LLM in the target languages increases the quality of generated data for downstream model performance substantially when compared to English demonstrations. This is the case in nearly all cases. However, the inclusion of English demonstration is also generally better for data quality compared to cases where only summarized label descriptions are used.

## 4.5 Analysis of LLM Revision Strategy

Revision consistently improves generation quality as demonstrated in the downstream performance comparison across different tasks, models and languages. Table 8 shows how many samples are rejected in the revision cycle by each of the models. For Llama models, we can see a clear pattern that a smaller version with 8 billion parameters rejects more samples than the model with 70 billion parameters. For instance, for the topic classification task, Llama3-8b rejected 21% of all generated samples while Llama3-70b rejected only 0.89%. Similarly, for the same task, Gemma3-4b rejected 24.07% of samples, and Gemma3-27b rejected only 9.64%. However, for sentiment analysis, both versions of Gemma reject approximately the same amount (around 7.5%), while Llama models retain the pattern where the larger version has a much lower rejection rate (only 1%).

Our manual analysis revealed that the rejected samples typically describe related concepts that are thematically relevant but do not fully describe the main intent (e.g. travel expenses and booking in Figure 2 are only indirectly related to travel). The revision also helps to filter out the cases that are ungrammatical, have incorrect output language, or are irrelevant to the intent, which is often the problem for low-resource languages.

Based on the rejection proportions across languages (see Table 8) we can conclude that generation quality differs substantially depending on the language. For example, Welsh samples have a high rate of rejection (up to 65.43% for the topic classification task with Gemma3-4b) while English samples have comparatively low rejection rates.

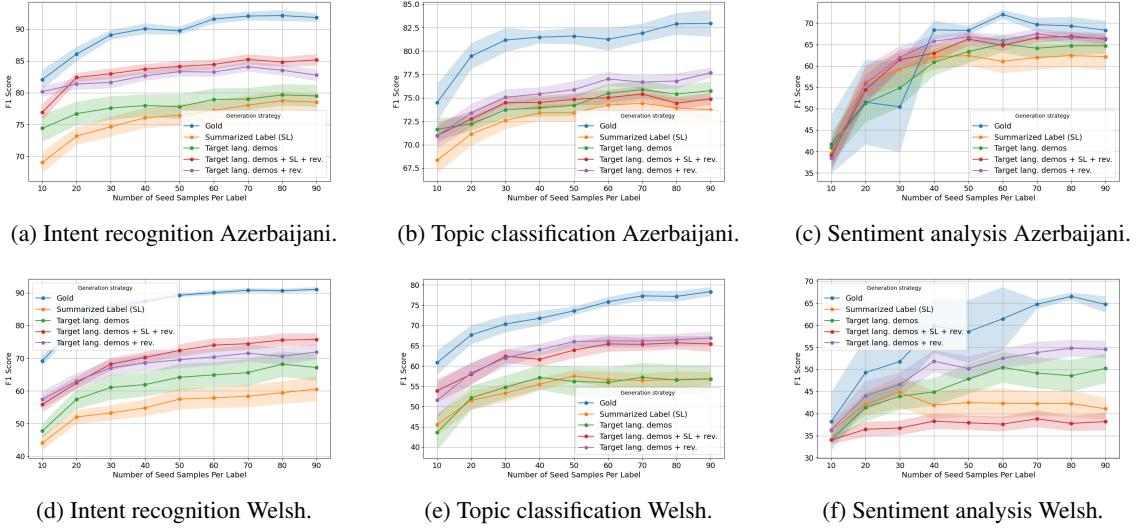


Figure 4: F1 mean performance with 95% Conf. Intervals for low-resource languages of Azerbaijani and Welsh for different numbers of seed samples per label used. We visualize the most prominent generation strategies and the gold data for comparison.

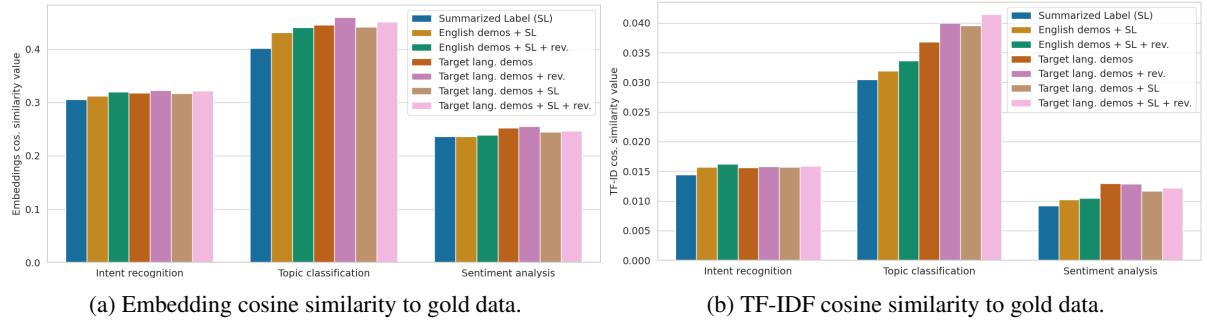


Figure 5: Similarity of data generated through various generation strategies to gold data for two different similarity metrics. The figures display results averaged across languages and models.

## 5 Additional Evaluation: Similarity of Generated Data to Gold Data

In order to get an idea of how close the generated data is semantically to the original data, we measure the average similarity of each generation strategy to gold data in Figure 5, where we distinguish between various tasks in our study. We compute two metrics: The cosine similarity of TF-IDF (Rajaraman and Ullman, 2011) and the cosine similarity of embeddings. For the embedding computation, we used a multilingual embedding model (Feng et al., 2022) and calculated similarity pairwise, using the samples with the same label, and then averaged the scores.

Notably, LLM revision increases the similarity of the generated data with regard to gold data in nearly all cases where revision is used. This is most notable for topic classification, while the difference

in similarity for intent recognition and sentiment analysis is quite low. Target language demonstrations also lead to generally higher similarity between the gold and the generated data.

Interestingly, although the similarity to the gold data between the different generation strategies for the intent recognition task does not vary much, some strategies still result in better performance than others, with target demonstrations and revision being the most beneficial, as demonstrated in Figure 3a. This indicates that, depending on the task, generated examples can be useful for the downstream task even when they are not very similar to the original gold samples, potentially even adding more diversity to the training data.

We note that the Pearson correlation coefficient between the F1 performance of individual methods and their cosine similarity to human data correlates positively, but only very weakly (0.111 on average

across tasks and languages).

To check our hypothesis that generated data add more diversity, we evaluated n-gram diversity of the gold and generated training samples (see A.13) and found that the diversity of the data generated with revision and demonstrations is on average +0.218 higher for the intent recognition task, and it is also higher for the generated sentiment analysis data (+0.581), although the metric is slightly lower for the topic classification (-0.279 compared to the gold data). This can be explained by the fact that the sentiment data includes the most diverse samples that were sourced from different datasets, varying in length and topic coverage. Also, data generated with Gemma demonstrates, on average, better diversity than Llama-generated data, especially with the larger version of Gemma that has 27b parameters.

## 6 Conclusion

In this work, we perform a comparative analysis of seven data generation strategies that combine summarised label descriptions, language of demonstrations, and LLM-based revision. We benchmark the models fine-tuned on the data generated with these strategies on three NLP tasks and eleven languages, including several low-resourced ones. Our experiments reveal that the best strategy, in general, is the usage of target language demonstrations with revision: revision consistently improves the quality of generated samples even in the low-resource scenario. Next, for the Llama models, target language demonstrations can be replaced with examples in English for comparable performance. We also found that while intent recognition tasks benefit from the inclusion of label descriptions, generating the data for topic classification works better without them.

As expected, larger models tend to generate better quality examples, and Llama models slightly outperform Gemma. However, even though data generated for higher-resource languages has better quality, using an optimal generation strategy may substantially reduce the performance gap. Our findings highlight the importance of careful selection of the data generation strategy, as we provide a recipe for creating generation prompts in a low-resource setting for different tasks and languages.

## Limitations

Due to the scope of this study and the large number of tested configurations, we experiment with only two model families with good multilingual capabilities: Gemma-3 and Llama-3. It is possible that some other models, e.g. Qwen-3, will perform better in the generation setting, at least for some of the languages. We also limited the study to only 11 languages due to resource constraints and the availability of evaluation data for the low-resource languages.

We could also consider all available labels in the MASSIVE dataset (currently subsampled to 10). However, we believe that such simplifications were necessary to avoid many conflating factors that can be potentially caused by the semantic overlaps in label descriptions.

We additionally did not consider using English demonstrations without revision. Target language demonstrations with revision consistently outperform English demonstrations with revision, and as such, we would expect similar performance for English demonstrations without revision compared to target language demonstrations without revision, making their usage in this study redundant.

Another limitation of this work is the lack of human annotations for checking the quality of the generated data. Involving human annotators was not feasible in the scope of the current project, given the amount of generated data for multiple languages, tasks, models, and generation configurations. Besides, many of the languages included in our study are very low-resourced (e.g. Welsh, Azerbaijani, Telugu), and it would be difficult to find native speakers to perform human evaluation for them. We consider a follow-up study with human annotators as future work, and in the current project, we opt for an automated evaluation using the downstream performance on the gold test data (that was human-annotated) as a performance indicator.

We recognise that the extent of potential data contamination is currently unknown because the data on which LLMs were trained is not fully disclosed, which might be reflected in the disparities between the LLMs' performance when generating data in different languages and domains. We believe that measuring this reliably is out-of-scope for our study, but we recognise this as a potential problem.

## Acknowledgments

This work was supported by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90254); by the German Federal Ministry of Research, Technology and Space (BMFTR) as part of the project TRAILS (01IW24005); by DisAI and AI-CODE, projects funded by European Union under the Horizon Europe, GA No. 101079164 and No. 101135437; by the European Union NextGenerationEU through the Recovery and Resilience Plan for Slovakia under the projects No. 09I01-03-V04-00059. and No. 09I03-03-V03-00020.

## References

Louis Abraham, Charles Arnal, and Antoine Marie. 2025. [Prompt selection matters: Enhancing text annotations for social sciences with large language models](#). *Preprint*, arXiv:2407.10645.

David Ifeoluwa Adelani, Hannah Liu, Xiaoyu Shen, Nikita Vassilyev, Jesujoba O. Alabi, Yanke Mao, Haonan Gao, and En-Shiun Annie Lee. 2024. [SIB-200: A simple, inclusive, and big evaluation dataset for topic classification in 200+ languages and dialects](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 226–245, St. Julian’s, Malta. Association for Computational Linguistics.

AI@Meta. 2024. [Llama 3 model card](#).

Jan Cegin, Branislav Pecher, Jakub Simko, Ivan Srba, Maria Bielikova, and Peter Brusilovsky. 2024a. [Effects of diversity incentives on sample diversity and downstream model performance in LLM-based text augmentation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13148–13171, Bangkok, Thailand. Association for Computational Linguistics.

Jan Cegin, Branislav Pecher, Jakub Simko, Ivan Srba, Maria Bielikova, and Peter Brusilovsky. 2024b. [Use random selection for now: Investigation of few-shot selection strategies in llm-based text augmentation for classification](#). *Preprint*, arXiv:2410.10756.

Jan Cegin, Jakub Simko, and Peter Brusilovsky. 2023. [ChatGPT to replace crowdsourcing of paraphrases for intent classification: Higher diversity and comparable model robustness](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1889–1905, Singapore. Association for Computational Linguistics.

Jan Cegin, Jakub Simko, and Peter Brusilovsky. 2025. [LLMs vs established text augmentation techniques for classification: When do the benefits outweigh the costs?](#) In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 10476–10496, Albuquerque, New Mexico. Association for Computational Linguistics.

Yi-Ling Chung, Aurora Cobo, and Pablo Serna. 2025. [Beyond translation: Llm-based data generation for multilingual fact-checking](#). *arXiv preprint arXiv:2502.15419*.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale](#). *CoRR*, abs/1911.02116.

Haixing Dai, Zhengliang Liu, Wenxiong Liao, Xiaoke Huang, Yihan Cao, Zihao Wu, Lin Zhao, Shaochen Xu, Wei Liu, Ninghao Liu, Sheng Li, Dajiang Zhu, Hongmin Cai, Lichao Sun, Quanzheng Li, Dinggang Shen, Tianming Liu, and Xiang Li. 2023. [Auggpt: Leveraging chatgpt for text data augmentation](#). *Preprint*, arXiv:2302.13007.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Chandra Kiran Reddy Evuru, Sreyan Ghosh, Sonal Kumar, S Ramaneswaran, Utkarsh Tyagi, and Dinesh Manocha. 2024. [Coda: Constrained generation based data augmentation for low-resource nlp](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3754–3769.

Luyang Fang, Gyeong-Geon Lee, and Xiaoming Zhai. 2023. [Using gpt-4 to augment unbalanced data for automatic scoring](#). *Preprint*, arXiv:2310.18365.

Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.

Steven Y. Feng, Varun Gangal, Jason Wei, Sarah Chandar, Soroush Vosoughi, Teruko Mitamura, and Eduard Hovy. 2021. [A survey of data augmentation approaches for NLP](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 968–988, Online. Association for Computational Linguistics.

Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron

Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha Britan, Wouter Leeuwis, Gokhan Tur, and Prem Natarajan. 2023. **MASSIVE: A 1M-example multilingual natural language understanding dataset with 51 typologically-diverse languages**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4277–4302, Toronto, Canada. Association for Computational Linguistics.

Sreyan Ghosh, Chandra Kiran Evuru, Sonal Kumar, S Ramaneswaran, S Sakshi, Utkarsh Tyagi, and Dinesh Manocha. 2023. Dale: Generative data augmentation for low-resource legal nlp. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Sentosa, Singapore.

Parker Glenn, Alolika Gon, Nikhil Kohli, Sihan Zha, Parag Pravin Dakle, and Preethi Raghavan. 2023. **Jetsons at the FinNLP-2023: Using synthetic data and transfer learning for multilingual ESG issue classification**. In *Proceedings of the Fifth Workshop on Financial Technology and Natural Language Processing and the Second Multimodal AI For Financial Forecasting*, pages 133–139, Macao. -.

Xu Gu, Xiaoliang Chen, Peng Lu, Zonggen Li, Yajun Du, and Xianyong Li. 2024. **Agcvt-prompt for sentiment classification: Automatically generating chain of thought and verbalizer in prompt learning**. *Eng. Appl. Artif. Intell.*, 132(C).

Daniil Gurgurov, Mareike Hartmann, and Simon Ostermann. 2024. **Adapting multilingual LLMs to low-resource languages with knowledge graphs via adapters**. In *Proceedings of the 1st Workshop on Knowledge Graphs and Large Language Models (KaLLM 2024)*, pages 63–74, Bangkok, Thailand. Association for Computational Linguistics.

Daniil Gurgurov, Rishu Kumar, and Simon Ostermann. 2025a. **GrEmLIn: A repository of green baseline embeddings for 87 low-resource languages injected with multilingual graph knowledge**. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 1204–1221, Albuquerque, New Mexico. Association for Computational Linguistics.

Daniil Gurgurov, Ivan Vykopal, Josef van Genabith, and Simon Ostermann. 2025b. Small models, big impact: Efficient corpus and graph-based adaptation of small multilingual language models for low-resource languages. *arXiv preprint arXiv:2502.10140*.

Alina Kramchaninova and Arne Defauw. 2022. **Synthetic data generation for multilingual domain-adaptable question answering systems**. In *Proceedings of the 23rd Annual Conference of the European Association for Machine Translation*, pages 151–160, Ghent, Belgium. European Association for Machine Translation.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

Xiaonan Li, Changtai Zhu, Linyang Li, Zhangyue Yin, Tianxiang Sun, and Xipeng Qiu. 2024. **LLatireval: LLM-verified retrieval for verifiable generation**. 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 5453–5471, Mexico City, Mexico. Association for Computational Linguistics.

Linlin Liu, Bosheng Ding, Lidong Bing, Shafiq Joty, Luo Si, and Chunyan Miao. 2021. **MulDA: A multilingual data augmentation framework for low-resource cross-lingual NER**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5834–5846, Online. Association for Computational Linguistics.

Qijiong Liu, Nuo Chen, Tetsuya Sakai, and Xiao-Ming Wu. 2024. **Once: Boosting content-based recommendation with both open- and closed-source large language models**. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, WSDM '24, page 452–461, New York, NY, USA. Association for Computing Machinery.

Sepideh Mollanorozy, Marc Tanti, and Malvina Nissim. 2023. **Cross-lingual transfer learning with Persian**. In *Proceedings of the 5th Workshop on Research in Computational Linguistic Typology and Multilingual NLP*, pages 89–95, Dubrovnik, Croatia. Association for Computational Linguistics.

Amani Namboori, Shivam Sadashiv Mangale, Andy Rosenbaum, and Saleh Soltan. 2023. **Gemquad: Generating multilingual question answering datasets from large language models using few shot learning**. In *NeurIPS 2023 Workshop on Synthetic Data Generation with Generative AI*.

Aytuğ Onan. 2023. **Srl-aco: A text augmentation framework based on semantic role labeling and ant colony optimization**. *Journal of King Saud University - Computer and Information Sciences*, 35(7):101611.

Frédéric Piedboeuf and Philippe Langlais. 2023. **Is ChatGPT the ultimate data augmentation algorithm?** In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15606–15615, Singapore. Association for Computational Linguistics.

Salsabila Zahrah Pranida, Rifo Ahmad Genadi, and Fajri Koto. 2025. **Synthetic data generation for culturally nuanced commonsense reasoning in low-resource languages**. *Preprint*, arXiv:2502.12932.

Anand Rajaraman and Jeffrey David Ullman. 2011. *Data Mining*, page 1–17. Cambridge University Press.

Gaurav Sahu, Pau Rodriguez, Issam Laradji, Parmida Atighehchian, David Vazquez, and Dzmitry Bahdanau. 2022. [Data augmentation for intent classification with off-the-shelf large language models](#). In *Proceedings of the 4th Workshop on NLP for Conversational AI*, pages 47–57, Dublin, Ireland. Association for Computational Linguistics.

Indira Sen, Dennis Assenmacher, Mattia Samory, Isabelle Augenstein, Wil Aalst, and Claudia Wagner. 2023. [People make better edits: Measuring the efficacy of LLM-generated counterfactually augmented data for harmful language detection](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10480–10504, Singapore. Association for Computational Linguistics.

Tran Ngoc Son, Nguyen Anh Tu, and Nguyen Minh Tri. 2025. [An efficient approach for machine translation on low-resource languages: A case study in vietnamese-chinese](#). *arXiv preprint arXiv:2501.19314*.

Gemma Team. 2025. [Gemma 3](#).

Solomon Ubani, Suleyman Olcay Polat, and Rodney Nielsen. 2023. [ZeroShotdataaug: Generating and augmenting training data with chatgpt](#). *Preprint, arXiv:2304.14334*.

Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. [Do llamas work in English? on the latent language of multilingual transformers](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15366–15394, Bangkok, Thailand. Association for Computational Linguistics.

Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu, Dacheng Tao, and Tianyi Zhou. 2024. [A survey on knowledge distillation of large language models](#). *arXiv preprint arXiv:2402.13116*.

Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyoung Park. 2021. [GPT3Mix: Leveraging large-scale language models for text augmentation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2225–2239, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Elena Zotova, Rodrigo Agerri, and German Rigau. 2021. [Semi-automatic generation of multilingual datasets for stance detection in twitter](#). *Expert Systems with Applications*, 170:114547.

## A Appendix

### A.1 Ethical Considerations

Based on a thorough ethical assessment performed on the basis of intra-institutional ethical guidelines and checklists tailored to the use of data and algorithms, we see no ethical concerns pertaining directly to the conduct of this research. Although the production of new data through LLMs bears several risks, such as the introduction of biases, the small size of the produced dataset, sufficient for experimentation, is, at the same time, insufficient for any major machine learning endeavours where such biases could be transferred.

We follow the license terms for all the models and datasets we used (such as the one required for the use of the Llama-3 model) – all models and datasets allow their use as part of the research.

### A.2 Language Abbreviations

Code	Language
az	Azerbaijani
cy	Welsh
de	German
en	English
he	Hebrew
id	Indonesian
ro	Romanian
sl	Slovenian
sw	Swahili
te	Telugu
th	Thai

Table 2: Language abbreviations.

### A.3 Computational Resources

All generation experiments were done using vllm (Kwon et al., 2023) for efficient inference. For a single language-model combination it takes approximately 2 hours to generate the data for intent recognition, 1 hour 20 minutes for topic classification and 23 minutes for sentiment classification. We use a single H100 GPU to generate the samples. Thus, for 11 languages, 4 models and 3 different tasks it takes around 157 hours (6.5 days) to generate all the data.

The fine-tuning experiments were performed on RTXA6000 GPU, and XLM-R fine-tuning for each setting with 10 different seeds takes around 120 minutes for intent recognition, 95 minutes for topic classification and 30 minutes for sentiment analysis. Note that all computations can be done in parallel.

### A.4 Intent Recognition Task: Examples

label	example
datetime_convert	<i>tell me the time in g. m. t. plus five</i>
alarm_query	<i>please list active alarms</i>
weather_query	<i>how's the weather like in beijing</i>
audio_volume_down	<i>you're too loud</i>
play_audiobook	<i>i want to start war and peace where i left off</i>
cooking_recipe	<i>what's needed to make pizza</i>
recommendation_movies	<i>name a rom com movie playing in and around new york theaters</i>
transport_ticket	<i>olly i need to get to bristol friday night can you book me a ticket please</i>
email_sendemail	<i>could you please gather a list of local restaurants and email them to my husband</i>
calendar_remove	<i>remove stand-up on friday at ten am</i>

Table 3: Labels from the MASSIVE dataset.

### A.5 Topic Classification Task: Examples

label	example
health	<i>Danielle Lantagne, a UN expert on the disease, stated the outbreak was likely caused by the peacekeepers.</i>
geography	<i>Large areas further north are quite sparsely populated, and some are nearly uninhabited wilderness.</i>
politics	<i>Soon after the outbreak of hostilities, Britain initiated a naval blockade of Germany.</i>
entertainment	<i>“She’s very cute and sings quite well, too,” he said according to a transcript of the news conference.</i>
science/technology	<i>The atom can be considered to be one of the fundamental building blocks of all matter.</i>
sports	<i>Massa is due to be out for at least the rest of the 2009 season.</i>
travel	<i>Over 60 cruise ships ply the Galapagos waters - ranging in size from 8 to 100 passengers.</i>

Table 4: Labels from the SIB-200 dataset.

### A.6 Downstream Fine-tuning of XLM-R

For the downstream evaluation we fine-tune XLM-R (Conneau et al., 2019) *FacebookAI/xlm-roberta-base* model for 50 epochs with a batch size 16 and employ early stopping with a patience of 5 epochs to prevent overfitting. We perform hyperparameter optimization to determine the optimal learning rate and set it to 1e-5, *AdamW* is used as an optimizer. Our dataset is balanced in both gold and generated settings and includes 100 samples per label. We

normalize all inputs by converting them to lower-case and removing punctuation. We fine-tune ten models per generation strategy for all language, task and model combinations, and report average F1 scores.

### A.7 Summarized Label Examples

This section shows the examples of summarized label descriptions (generated with Llama3-70b) for each task. Table 5 demonstrates generated descriptions for the intent recognition task, Table 6 for topic classification, and Table 7 for sentiment analysis.

### A.8 Prompt Templates

This section provides the prompt templates for sample generation (Figure 6), revision (Figure 7), and summarized label generation (Figure 8).

### A.9 Samples Rejected by Revision

Table 8 demonstrates the proportion of rejected samples in the revision setting.

### A.10 Model Dimension and Generation Settings

Figure 9 illustrates the average difference in performance when fine-tuning the downstream model on the generated data vs. gold data fine-tuning. The results are aggregated across all languages and tasks.

### A.11 Language Dimension and Generation Settings

Figure 10 illustrates the average difference in performance when fine-tuning the downstream model on the generated data vs. gold data fine-tuning. The results are aggregated across all models and tasks.

### A.12 Fine-tuning Results per Task, Model, and Language

This section demonstrates the fine-tuning results for all tasks and languages for Llama3-8b in Table 9, Gemma3-4b in Table 10, Gemma3-27b in Table 11, and Llama3-70b in Table 12.

### A.13 Data Diversity

Table 13 shows the results of computing the n-gram diversity for the original gold and the data generated with the target language demonstrations and revision. The results are aggregated across all languages.

intent	description
alarm_query	This intent involves querying or checking the status of existing alarms, including their timing, scheduling, and active status, to stay informed and manage one's alarm settings.
audio_volume_down	This intent involves requesting to reduce or lower the volume of an audio device, such as a speaker, to a softer level.
calendar_remove	This intent involves requesting to delete, remove, or cancel a specific event or all events from a calendar.
cooking_recipe	This intent involves seeking instructions or information about preparing and cooking various dishes, including specific ingredients, cooking times, temperatures, and methods.
datetime_convert	This intent involves converting or comparing time zones, including calculating time differences between specific locations or zones, or adjusting a given time to a different time zone.
email_sendemail	This intent involves requesting to send an email to a specific recipient, often with a mention of the content or purpose of the email, such as sharing information, making a request, or inquiring about something.
play_audiobook	This intent involves requesting to play, resume, or restart an audiobook, often specifying the title or author, and may include additional actions like going back to a previous point in the audiobook.
recommendation_movies	This intent involves asking for movie recommendations or inquiring about current movie showtimes and ratings, often with a focus on the user's location or specific preferences.
transport_ticket	This intent involves booking, purchasing, or searching for train tickets, often specifying details such as destination, travel date, and payment method.
weather_query	This intent involves asking about the current or future weather conditions, including temperature, precipitation, and other weather-related information, often specifying a particular location or date.

Table 5: Summarized label (intent) examples for the intent recognition task.

intent	description
entertainment	This intent involves discussing leisure activities, hobbies, and forms of entertainment, such as boating, music, comedy, and travel, as well as mentioning celebrities, events, and notable achievements in the entertainment industry.
geography	This intent involves describing or providing information about geographical locations, features, and formations, including cities, islands, mountains, and landforms, as well as their characteristics, sizes, and relationships to each other.
health	This intent involves discussing various aspects of health, including infectious and contagious diseases, their causes, symptoms, and prevention methods, as well as medical treatments, diagnostic tools, and personal health care practices.
politics	This intent involves discussions, statements, and actions related to government policies, political leaders, elections, and legislation, often referencing specific events, decisions, or opinions that shape or are shaped by political institutions and processes.
science/technology	This intent involves discussing or explaining scientific concepts, theories, and discoveries in various fields such as astronomy, medicine, biology, and technology, often providing historical context and details of experiments or research that support these concepts.
sports	This intent involves discussing various sports and athletic activities, including running, hockey, cross-country, skiing, fencing, golf, and others, as well as mentioning specific athletes, teams, and competitions.
travel	This intent involves planning, preparing for, or experiencing various modes of transportation, trips, and travel-related activities, including road travel, air travel, boat trips, and cruises, as well as considering factors that may impact travel, such as weather, road conditions, and cultural differences.

Table 6: Summarized label (intent) examples for the topic classification task.

intent	description
negative	This intent involves expressing negative sentiments, complaints, or criticisms about various topics, including work, relationships, social issues, and personal struggles. The examples convey frustration, disappointment, and even anger towards certain situations or individuals.
positive	This intent involves expressing positive sentiments, gratitude, and appreciation for various services, events, and announcements. The intent is to share good news, express joy, and acknowledge the efforts of individuals or organizations, often with a tone of optimism and enthusiasm.

Table 7: Summarized label (intent) examples for the sentiment classification task.

```
messages = [
{
    "role": "system",
    "content": f"You are an excellent text generator and can generate representative text samples for the given class in {lang_name}.",
},
{
    "role": "user", "content": f"You are required to produce {num_samples_to_generate} examples in {lang_name} that can have the label: {class_name} {added_explanation}. Generate {num_samples_to_generate} examples for the label {class_name}. Each example should be on a new line. Do not generate anything that cannot be classified as {class_name} and do not repeat the instruction. Generated examples for label {class_name}:"
},
]
```

Figure 6: Generation prompt.

```
messages = [
{
    "role": "system",
    "content": f"You are an excellent classifier and can reason whether a given sample in {lang_name} belongs to the class {class_name} or not.",
},
{
    "role": "user",
    "content": f"Decide whether the following example belongs to the class {class_name} which means {class_description}. Answer yes if it belongs and represents a good sample (grammatically correct and complete) and no if it does not. Answer no if the example is not in {lang_name}. Explain your answer in a concise way after generating yes or no. Input: {new_demo} Answer:",
},
]
```

Figure 7: Revision prompt.

```
messages = [
{
    "role": "system",
    "content": "You are an excellent summarizer and generate concise descriptions for different intents given some examples.",
},
{
    "role": "user",
    "content": f"Your task is to generate a short description of the intent {intent} based on the following examples from the dataset that share the intent {intent}: {examples} Your description should always start with 'This intent involves'. Description:",
},
]
```

Figure 8: Summarized label generation prompt.

	Topic Classification				Sentiment Analysis				Intent Recognition			
	Gemma3-4b	Gemma3-27b	Llama3-8b	Llama3-70b	Gemma3-4b	Gemma3-27b	Llama3-8b	Llama3-70b	Gemma3-4b	Gemma3-27b	Llama3-8b	Llama3-70b
Azerbaijani	34.29	12.86	40.71	1.14	14.50	9.50	46.50	1.50	22.50	28.80	43.60	4.30
Welsh	65.43	37.29	55.43	2.29	53.00	55.50	55.00	0.00	65.30	66.00	62.10	2.10
German	17.29	8.86	11.57	0.71	0.50	1.50	3.50	0.50	9.20	12.80	22.20	2.70
English	13.14	3.29	4.43	0.57	0.00	0.00	0.00	0.00	10.20	22.00	26.10	8.90
Hebrew	21.00	3.14	12.00	0.00	1.00	0.00	24.00	0.50	8.20	8.10	22.30	0.70
Indonesian	15.57	3.86	14.29	1.29	0.00	2.00	6.50	0.50	11.70	17.30	24.10	3.70
Romanian	15.71	2.86	9.71	0.71	0.00	3.00	8.50	0.50	6.80	17.40	22.00	2.80
Slovenian	15.86	6.29	16.86	0.57	0.00	2.50	26.50	3.50	21.30	21.20	20.00	4.10
Swahili	29.71	13.86	32.86	1.71	9.00	4.00	50.50	1.50	40.60	29.40	43.60	2.90
Telugu	17.43	11.00	9.43	0.14	1.50	5.00	11.00	2.00	8.90	13.50	2.90	0.30
Thai	19.29	2.71	24.00	0.71	1.00	0.00	20.00	1.00	11.20	9.80	31.50	2.20
total %	<b>24.07</b>	9.64	21.03	0.89	7.32	7.55	<b>22.91</b>	1.05	19.63	22.39	<b>29.13</b>	3.15

Table 8: Proportion of rejected samples in the revision generation setting with the intent description and target language demonstrations. Numbers in bold indicate the highest rejection rate per task.

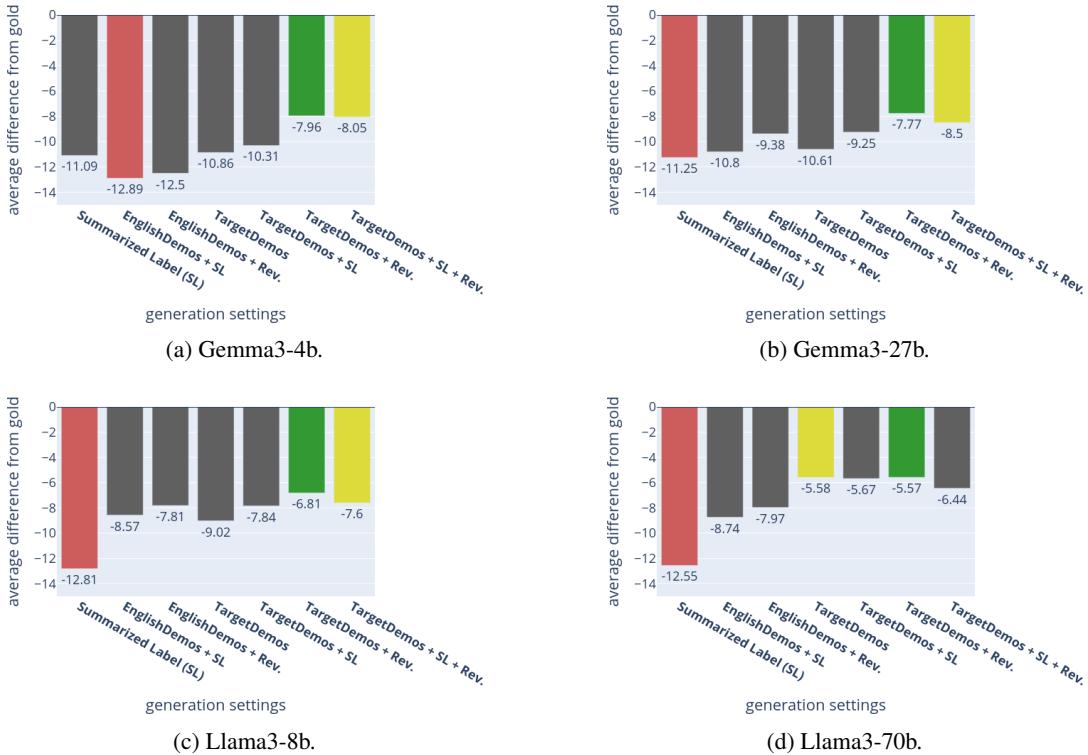


Figure 9: Average difference in performance on the generated data compared to the gold samples across the model dimension. Lower bars indicate better performance. Green bars show the best-performing generation setting across all languages and tasks for each model, and yellow bars show the second-best generation setting.

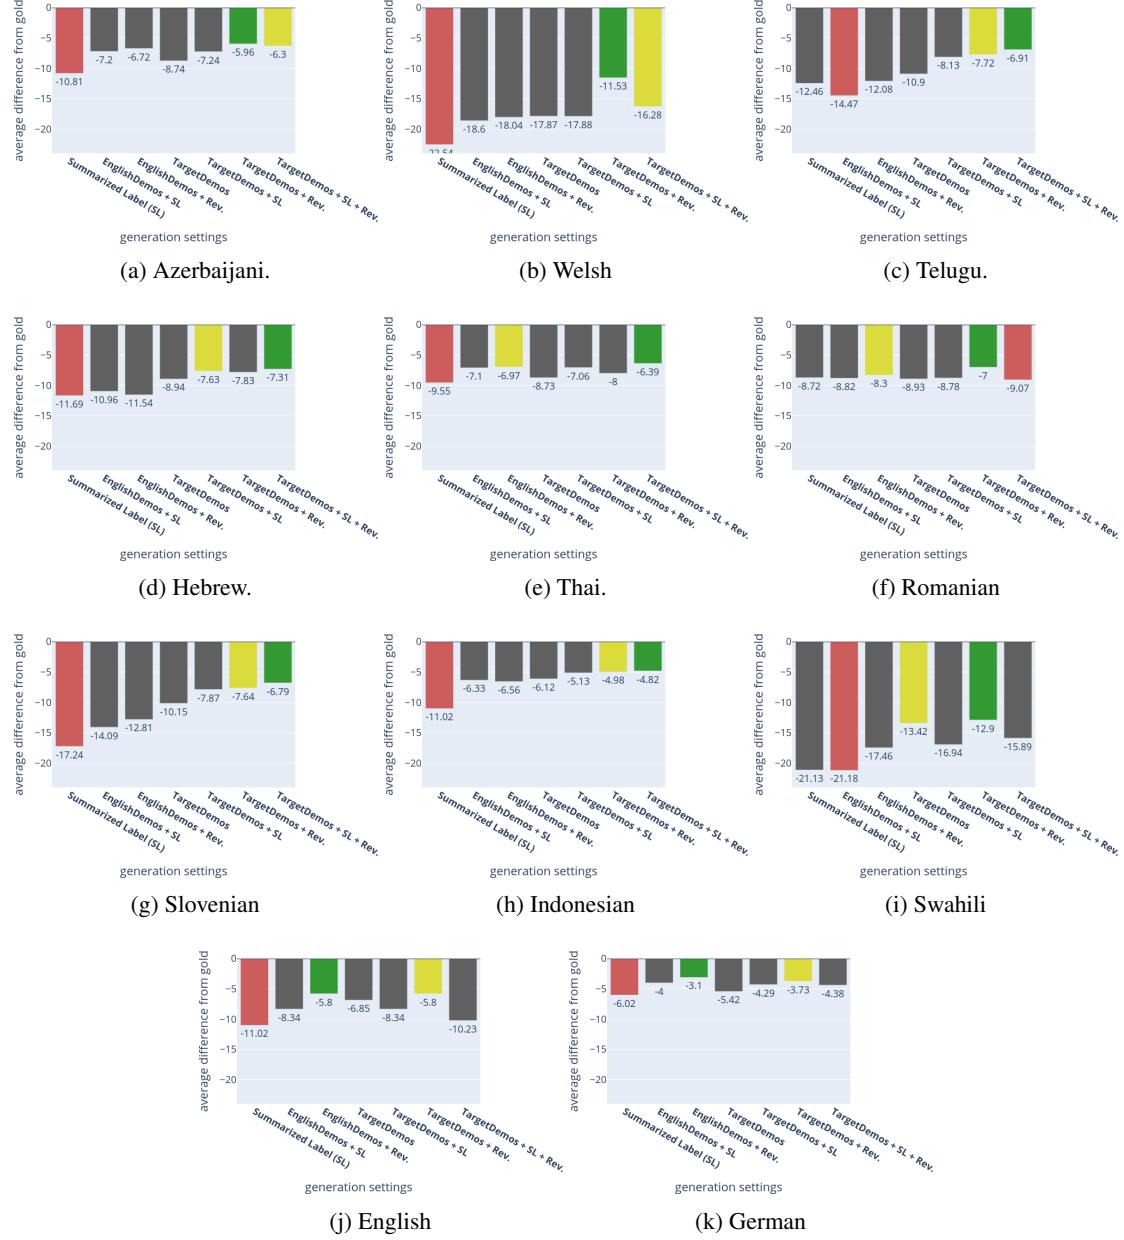


Figure 10: Average difference in performance on the generated data compared to the gold (original) samples across the language dimension. Lower bars indicate better performance. Green bars shows the best-performing generation setting across all models and tasks for each language, and yellow bars show the second best generation setting.

<b>Intent Recognition</b>	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	91.62	91.46	94.38	95.32	92.17	92.78	94.27	93.67	89.75	90.62	94.27	92.76
Summarized Label (SL)	73.80	59.94	90.01	<b>94.17</b>	70.35	90.46	90.11	87.30	60.62	75.21	85.40	79.76
EnglishDemos + SL	82.12	64.48	<u>92.03</u>	92.75	81.52	90.58	<b>92.25</b>	89.30	64.28	77.71	86.23	83.02
EnglishDemos + Rev.	83.56	70.19	92.02	<u>93.25</u>	81.77	90.96	<u>91.56</u>	84.96	65.80	81.93	86.29	83.85
TargetDemos	<u>85.02</u>	76.05	90.29	91.31	<b>87.75</b>	90.83	89.97	89.67	<u>79.41</u>	82.85	<u>89.49</u>	86.60
TargetDemos + SL	<b>85.09</b>	<b>77.65</b>	91.70	92.75	87.16	89.59	90.79	<b>91.68</b>	78.96	79.59	87.28	86.57
TargetDemos + Rev.	83.61	70.76	91.03	93.04	87.37	<u>91.63</u>	91.27	90.08	78.60	<b>84.48</b>	<b>90.36</b>	<u>86.57</u>
TargetDemos + SL + Rev.	85.00	<u>77.07</u>	<b>92.74</b>	93.25	<u>87.58</u>	<b>91.85</b>	90.92	<u>90.71</u>	<b>80.23</b>	<u>83.30</u>	87.44	<b>87.28</b>
<b>Topic Classification</b>	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	84.85	77.03	84.44	89.99	83.65	87.77	86.53	86.04	76.64	80.29	86.94	84.01
Summarized Label (SL)	72.80	61.72	71.77	70.68	73.46	73.44	72.18	71.52	63.41	69.45	71.91	70.21
EnglishDemos + SL	73.33	62.88	<b>78.19</b>	<b>79.65</b>	<b>76.74</b>	79.83	75.17	75.95	62.59	66.72	75.63	73.34
EnglishDemos + Rev.	<b>77.55</b>	67.85	75.16	77.52	<u>75.55</u>	<b>80.83</b>	<u>76.05</u>	80.22	<b>65.71</b>	70.30	<b>77.66</b>	<u>74.95</u>
TargetDemos	73.53	64.94	75.17	77.19	72.95	78.62	74.92	78.41	55.83	<b>76.61</b>	72.45	72.78
TargetDemos + SL	68.60	66.48	<u>78.08</u>	<u>79.65</u>	73.07	78.87	74.81	<u>81.25</u>	<u>65.58</u>	73.09	76.28	74.16
TargetDemos + Rev.	<u>74.89</u>	<b>70.46</b>	77.58	<u>79.19</u>	73.21	<u>79.88</u>	<b>77.02</b>	<b>81.45</b>	60.77	<u>74.86</u>	<u>76.41</u>	<b>75.06</b>
TargetDemos + SL + Rev.	72.97	<u>69.81</u>	77.30	77.52	73.69	76.53	72.20	80.10	62.03	72.61	75.02	73.62
<b>Sentiment Analysis</b>	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	71.69	58.77	65.84	80.97	82.23	90.39	90.48	83.40	75.07	82.51	80.14	78.32
Summarized Label (SL)	46.59	47.16	66.83	67.99	72.45	74.54	82.41	50.69	42.02	<b>79.61</b>	69.95	63.66
EnglishDemos + SL	66.16	40.99	68.95	<u>75.57</u>	72.86	86.43	86.05	<u>74.81</u>	<u>60.90</u>	<u>73.91</u>	<b>78.30</b>	<u>71.36</u>
EnglishDemos + Rev.	<b>71.33</b>	35.20	<b>73.08</b>	70.48	<u>75.28</u>	<u>88.29</u>	83.58	67.66	<b>66.22</b>	66.16	<b>79.33</b>	70.60
TargetDemos	54.20	<u>52.83</u>	67.65	<b>78.72</b>	66.97	83.06	82.54	55.41	57.97	61.80	74.74	66.90
TargetDemos + SL	64.17	43.20	<u>68.98</u>	75.57	<b>75.59</b>	<b>88.90</b>	82.64	71.75	53.51	63.98	72.76	69.19
TargetDemos + Rev.	66.94	<b>53.79</b>	65.27	<u>77.95</u>	69.74	84.76	<b>87.11</b>	74.40	60.55	70.14	76.56	<b>71.57</b>
TargetDemos + SL + Rev.	67.23	34.06	61.59	70.48	73.83	85.99	<u>86.48</u>	<b>76.28</b>	58.53	67.95	77.78	69.11

Table 9: Fine-tuning results for the gold and artificial data generated by Llama3-8b on the three tasks. For each configuration we fine-tune ten downstream XLM-R models and report average F1 scores. We bold the best generation strategy in each column and underline the second best. See A.2 for language abbreviations.

<b>Intent Recognition</b>	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	91.62	91.46	94.38	95.32	92.17	92.78	94.27	93.67	89.75	90.62	94.27	92.76
Summarized Label (SL)	77.36	52.51	84.09	92.47	82.58	89.02	88.85	89.25	76.49	83.59	85.34	81.96
EnglishDemos + SL	76.39	61.24	91.19	92.89	85.12	91.17	89.41	84.92	63.73	79.21	86.30	81.96
EnglishDemos + Rev.	79.84	47.33	<b>92.90</b>	<b>92.94</b>	80.30	<u>91.27</u>	86.12	82.76	66.21	77.83	87.92	80.49
TargetDemos	76.40	49.11	90.33	89.05	<u>89.39</u>	87.86	88.82	87.83	72.52	82.36	89.36	82.09
TargetDemos + SL	<b>85.55</b>	<u>65.71</u>	91.07	92.89	<b>89.41</b>	<b>91.91</b>	<b>91.54</b>	<u>90.52</u>	<u>77.47</u>	<u>87.54</u>	<b>90.63</b>	<u>86.75</u>
TargetDemos + Rev.	82.22	61.56	89.85	91.25	88.67	89.72	90.21	88.16	76.25	84.89	<u>89.57</u>	84.76
TargetDemos + SL + Rev.	84.65	<b>71.88</b>	<u>91.79</u>	<u>92.94</u>	89.26	90.96	<u>90.79</u>	<b>91.58</b>	<b>79.64</b>	<b>88.01</b>	89.08	<b>87.33</b>
<b>Topic Classification</b>	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	81.94	78.61	84.91	89.99	83.98	85.33	86.38	85.83	76.71	80.31	87.51	83.77
Summarized Label (SL)	73.61	49.97	71.56	75.50	70.39	70.58	<b>75.20</b>	72.12	53.52	65.63	<u>77.26</u>	68.67
EnglishDemos + SL	<b>79.39</b>	45.36	<b>74.82</b>	76.30	63.02	70.47	74.76	<u>75.23</u>	59.42	55.90	<b>78.06</b>	68.43
EnglishDemos + Rev.	76.12	54.77	<u>74.22</u>	76.65	66.58	<u>77.98</u>	74.24	<b>75.40</b>	<b>61.98</b>	58.16	76.47	70.23
TargetDemos	76.89	44.20	70.67	<b>80.15</b>	70.57	76.92	72.74	74.64	56.27	66.00	72.76	69.26
TargetDemos + SL	69.00	38.52	70.93	76.30	<b>70.95</b>	73.20	72.49	73.26	60.34	68.53	72.73	67.84
TargetDemos + Rev.	<u>79.22</u>	<b>64.20</b>	73.95	<u>79.72</u>	<u>70.79</u>	<b>79.22</b>	74.14	73.73	57.38	<b>73.33</b>	71.72	<b>72.49</b>
TargetDemos + SL + Rev.	75.20	<u>57.67</u>	72.23	76.65	70.02	77.08	<u>74.83</u>	73.49	<u>61.49</u>	71.53	74.17	71.31
<b>Sentiment Analysis</b>	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	71.69	58.77	65.84	80.97	82.23	90.39	90.48	83.40	75.07	82.51	80.14	78.32
Summarized Label (SL)	<b>69.86</b>	33.88	64.50	68.52	<b>77.99</b>	<b>88.19</b>	83.42	68.20	56.89	70.00	69.12	68.23
EnglishDemos + SL	68.09	33.72	60.40	69.08	62.12	<u>77.55</u>	81.52	70.38	34.71	61.18	<u>76.95</u>	63.24
EnglishDemos + Rev.	65.69	33.87	62.91	66.17	53.20	77.82	83.88	60.51	49.65	74.13	74.34	63.83
TargetDemos	67.98	33.70	63.70	<u>73.19</u>	69.23	79.08	<u>84.67</u>	<u>78.02</u>	<u>61.31</u>	72.83	72.44	68.74
TargetDemos + SL	<b>70.77</b>	33.63	64.27	69.08	<u>77.84</u>	80.34	76.22	77.98	30.57	78.56	75.27	66.78
TargetDemos + Rev.	66.27	<b>49.17</b>	<b>68.97</b>	<b>77.82</b>	74.76	77.51	<b>85.62</b>	76.77	<b>66.30</b>	<b>80.30</b>	69.97	<b>72.13</b>
TargetDemos + SL + Rev.	69.71	<u>45.02</u>	<u>66.36</u>	66.17	77.45	<u>85.20</u>	84.28	<b>79.42</b>	31.66	78.76	<b>78.08</b>	69.28

Table 10: Fine-tuning results for the gold and artificial data generated by Gemma3-4b on the three tasks. For each configuration we fine-tune ten downstream XLM-R models and report average F1 scores. We bold the best generation strategy in each column and underline the second best. See A.2 for language abbreviations.

<b>Intent Recognition</b>	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	91.62	91.46	94.38	95.32	92.17	92.78	94.27	93.67	89.75	90.62	94.27	92.76
Summarized Label (SL)	76.28	74.84	87.96	92.42	85.06	89.82	88.81	88.05	80.05	84.88	87.98	85.10
EnglishDemos + SL	86.66	74.02	92.51	<b>93.02</b>	85.58	92.91	87.14	81.71	81.58	77.39	86.45	85.36
EnglishDemos + Rev.	87.15	77.91	92.64	92.93	<b>91.32</b>	<u>93.28</u>	<b>92.79</b>	<u>91.69</u>	<b>82.38</b>	87.30	85.80	88.65
TargetDemos	73.97	66.66	88.79	89.56	89.78	90.52	91.24	<b>91.90</b>	80.61	83.07	85.68	84.71
TargetDemos + SL	<u>87.17</u>	70.97	91.52	<u>93.02</u>	88.97	91.79	<u>91.49</u>	90.55	79.51	<b>88.38</b>	<u>88.95</u>	87.48
TargetDemos + Rev.	86.08	76.14	91.60	92.14	<u>90.29</u>	91.15	91.46	91.64	81.28	87.18	86.99	87.81
TargetDemos + SL + Rev.	<b>88.81</b>	<b>79.35</b>	<b>93.16</b>	92.93	88.80	<b>93.65</b>	90.15	88.48	<u>82.11</u>	<u>87.84</u>	<b>90.19</b>	<b>88.68</b>
<b>Topic Classification</b>	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	81.94	78.61	84.91	89.99	83.98	85.33	86.38	85.83	76.71	80.31	87.51	83.77
Summarized Label (SL)	73.92	59.52	72.63	72.08	65.63	69.27	<u>73.39</u>	71.80	62.16	64.99	72.36	68.89
EnglishDemos + SL	<u>77.12</u>	<b>64.35</b>	73.43	75.12	<b>71.09</b>	74.44	69.20	70.82	62.26	<b>70.28</b>	75.66	71.25
EnglishDemos + Rev.	76.50	<b>65.93</b>	73.81	74.49	<u>69.92</u>	72.05	72.03	<b>74.83</b>	<b>67.28</b>	<u>70.22</u>	<b>76.62</b>	<b>72.15</b>
TargetDemos	75.51	49.63	71.54	<u>76.93</u>	68.93	74.21	69.17	73.25	63.20	69.78	73.79	69.63
TargetDemos + SL	73.64	57.87	74.30	75.12	65.88	74.55	70.11	74.57	58.53	66.90	76.15	69.78
TargetDemos + Rev.	<b>77.36</b>	62.55	<u>74.69</u>	<b>77.72</b>	65.20	<b>75.49</b>	<b>75.80</b>	<u>74.70</u>	<u>67.22</u>	67.21	72.57	<u>71.86</u>
TargetDemos + SL + Rev.	74.52	62.14	<b>75.47</b>	74.49	68.34	<u>75.19</u>	70.99	71.86	59.22	69.33	<b>76.78</b>	70.76
<b>Sentiment Analysis</b>	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	71.69	58.77	65.84	80.97	82.23	90.39	90.48	83.40	75.07	82.51	80.14	78.32
Summarized Label (SL)	63.46	35.38	66.55	65.80	<b>78.04</b>	71.52	<b>81.42</b>	41.84	53.19	69.79	<b>75.28</b>	63.84
EnglishDemos + SL	<b>68.44</b>	35.05	62.93	62.39	72.65	83.69	78.81	46.55	33.32	70.32	74.15	62.57
EnglishDemos + Rev.	<u>65.44</u>	36.91	<u>67.58</u>	60.59	68.62	71.72	78.98	58.95	40.87	65.87	71.00	62.41
TargetDemos	63.33	<u>46.73</u>	<b>68.15</b>	<u>75.95</u>	73.87	87.41	72.86	58.32	<b>61.71</b>	53.56	69.66	66.51
TargetDemos + SL	61.89	33.87	65.02	62.39	78.01	<u>89.73</u>	<u>80.72</u>	65.70	49.77	<u>70.90</u>	<b>74.50</b>	66.59
TargetDemos + Rev.	63.34	<b>60.60</b>	65.99	<b>77.92</b>	<u>78.02</u>	<b>90.13</b>	74.99	73.78	<u>53.35</u>	62.75	71.14	<b>70.18</b>
TargetDemos + SL + Rev.	62.41	34.92	63.49	60.59	73.55	89.39	69.89	<b>75.20</b>	53.22	<b>75.60</b>	72.49	66.43

Table 11: Fine-tuning results for the gold and artificial data generated by Gemma3-27b on the three tasks. For each configuration we fine-tune ten downstream XLM-R models and report average F1 scores. We bold the best generation strategy in each column and underline the second best. See A.2 for language abbreviations.

Intent Recognition	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	91.62	91.46	94.38	95.32	92.17	92.78	94.27	93.67	89.75	90.62	94.27	92.76
Summarized Label (SL)	85.10	65.58	90.19	90.61	76.16	90.69	88.28	89.27	70.23	83.71	86.66	83.31
EnglishDemos + SL	83.82	76.08	<u>92.60</u>	92.50	81.88	90.56	88.08	88.02	74.64	84.32	86.05	85.32
EnglishDemos + Rev.	84.33	76.31	<u>92.02</u>	<b>92.99</b>	83.02	91.41	91.55	89.31	73.10	<u>85.42</u>	87.30	86.07
TargetDemos	84.24	<b>81.98</b>	91.97	92.26	<b>88.76</b>	<u>91.65</u>	90.40	88.13	80.03	<u>85.16</u>	<u>88.79</u>	87.58
TargetDemos + SL	85.00	79.89	<b>92.79</b>	92.50	88.65	90.79	<b>92.52</b>	<b>90.06</b>	<b>80.46</b>	84.94	88.20	<b>87.80</b>
TargetDemos + Rev.	<b>86.09</b>	79.66	91.53	92.96	88.46	<b>91.74</b>	<u>91.78</u>	<u>89.77</u>	<u>80.38</u>	<b>86.06</b>	86.92	<u>87.76</u>
TargetDemos + SL + Rev.	85.06	<u>80.28</u>	92.56	<u>92.99</u>	<u>88.74</u>	90.92	90.66	89.03	79.77	85.11	<b>89.40</b>	87.68
Topic Classification	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	84.85	77.03	84.44	89.99	83.65	87.77	86.53	86.04	76.64	80.29	86.94	84.01
Summarized Label (SL)	77.69	56.36	75.68	73.28	65.97	71.44	72.02	74.87	60.15	70.45	75.28	70.29
EnglishDemos + SL	70.73	70.75	73.25	<b>81.92</b>	<u>76.05</u>	<u>79.52</u>	73.56	78.85	68.80	72.03	77.71	74.83
EnglishDemos + Rev.	75.84	69.39	<b>78.27</b>	79.30	75.88	78.56	<u>76.94</u>	79.23	<b>70.71</b>	73.40	<b>80.70</b>	76.20
TargetDemos	<b>79.55</b>	<b>71.75</b>	75.51	80.03	75.90	78.88	<b>78.49</b>	<b>81.55</b>	68.93	73.33	76.96	<b>76.44</b>
TargetDemos + SL	75.63	70.47	76.33	<u>81.92</u>	75.37	78.63	76.70	79.13	68.78	<b>74.75</b>	79.13	76.08
TargetDemos + Rev.	<u>77.74</u>	<u>70.86</u>	<u>77.25</u>	80.78	76.04	<b>79.63</b>	75.79	<u>80.89</u>	67.98	<u>73.61</u>	79.66	<u>76.39</u>
TargetDemos + SL + Rev.	76.26	70.05	77.20	79.30	<b>76.89</b>	77.58	76.44	79.46	<u>69.47</u>	73.10	<u>80.23</u>	76.00
Sentiment Analysis	az	cy	de	en	he	id	ro	sl	sw	te	th	avg
Gold Data	71.69	58.77	65.84	80.97	82.23	90.39	90.48	83.40	75.07	82.51	80.14	78.32
Summarized Label (SL)	66.72	44.79	65.65	69.42	74.56	67.70	<u>84.06</u>	40.22	33.70	46.94	75.42	60.83
EnglishDemos + SL	68.15	<u>60.09</u>	<b>71.34</b>	73.90	72.78	85.70	82.94	46.36	45.58	51.12	<u>79.82</u>	67.07
EnglishDemos + Rev.	62.86	60.09	<u>67.83</u>	65.04	72.99	85.96	77.42	52.71	46.56	58.04	79.49	66.27
TargetDemos	71.32	<b>60.16</b>	60.84	<b>78.62</b>	71.44	86.41	81.83	73.14	<b>67.08</b>	75.51	75.72	<b>72.92</b>
TargetDemos + SL	<b>73.43</b>	59.41	63.21	73.90	70.45	<b>89.00</b>	79.47	71.11	59.17	<b>79.00</b>	<b>79.99</b>	72.56
TargetDemos + Rev.	<u>71.52</u>	54.08	67.13	<u>75.03</u>	<u>76.32</u>	<u>88.30</u>	<b>85.62</b>	64.93	<u>61.18</u>	76.22	78.72	<u>72.64</u>
TargetDemos + SL + Rev.	69.40	34.58	63.15	65.04	<b>77.02</b>	86.69	78.29	<b>74.88</b>	57.94	<u>77.64</u>	79.21	69.44

Table 12: Fine-tuning results for the gold and artificial data generated by Llama3-70b on the three tasks. For each configuration we fine-tune ten downstream XLM-R models and report average F1 scores. We bold the best generation strategy in each column and underline the second best. See A.2 for language abbreviations.

	Gold Data	Gemma3-4b	Llama3-8b	Gemma3-27b	Llama3-70b
Intent Recognition	2.475	2.913	2.414	<b>3.033</b>	2.412
Topic Classification	<b>3.377</b>	3.242	2.970	3.292	2.890
Sentiment Analysis	2.614	3.376	3.045	<b>3.437</b>	2.921

Table 13: N-gram diversity of the gold and generated data aggregated across all languages. The generation strategy includes demonstrations in the target language and revision. Numbers in bold indicate the largest diversity per task.