

# Better as Generators Than Classifiers: Leveraging LLMs and Synthetic Data for Low-Resource Multilingual Classification

Branislav Pecher<sup>†</sup>, Jan Cegin<sup>†</sup>, Robert Belanec<sup>♣†</sup>,  
Ivan Srba<sup>†</sup>, Jakub Simko<sup>†</sup>, Maria Bielikova<sup>†</sup>

<sup>†</sup> Kempelen Institute of Intelligent Technologies, Bratislava, Slovakia

<sup>♣</sup> Faculty of Information Technology, Brno University of Technology, Brno, Czechia  
{name.surname}@kinit.sk

## Abstract

Large Language Models (LLMs) have demonstrated remarkable multilingual capabilities, making them promising tools in both high- and low-resource languages. One particularly valuable use case is generating synthetic samples that can be used to train smaller models in low-resource scenarios where human-labelled data is scarce. In this work, we investigate whether these synthetic data generation capabilities can serve as a form of distillation, producing smaller models that perform on par with or even better than massive LLMs across languages and tasks. To this end, we use a state-of-the-art multilingual LLM to generate synthetic datasets covering 11 languages and 4 classification tasks. These datasets are then used to train smaller models via fine-tuning or instruction tuning, or as synthetic in-context examples for compact LLMs. Our experiments show that even small amounts of synthetic data enable smaller models to outperform the large generator itself, particularly in low-resource languages. Overall, the results suggest that LLMs are best utilised as generators (teachers) rather than classifiers, producing data that empowers smaller and more efficient multilingual models.

## 1 Introduction

Large Language Models (LLMs) have emerged as powerful generators of synthetic data, enabling the training of smaller, efficient downstream models across a wide range of tasks. It was successfully applied to different tasks, often rivalling or even surpassing traditional augmentation methods (Piedboeuf and Langlais, 2023; Onan, 2023; Cegin et al., 2023; Sen et al., 2023). Importantly, recent work has shown that LLMs are capable of producing high-quality synthetic data for low-resource languages (Anikina et al., 2025).

Besides generation, the LLMs are frequently used directly as classifiers through zero or few-shot prompting (Brown et al., 2020; Sun et al., 2023;

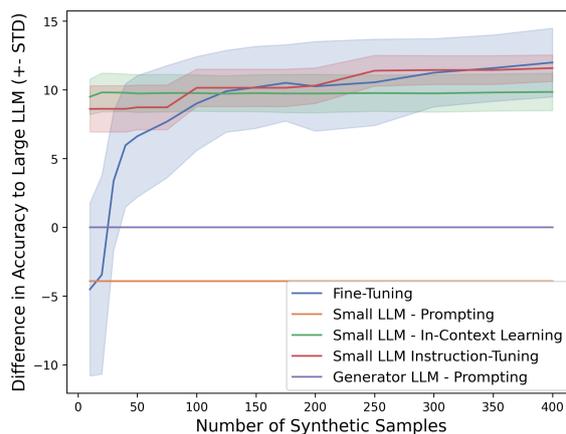


Figure 1: Smaller models trained with synthetic samples often outperform the large language model used for generating such samples. The results obtained by 5 models are aggregated over 11 languages and 4 tasks. The small LLM results are aggregated over 3 models.

Dong et al., 2022; Pecher et al., 2024b). A large body of work has examined their performance in this role, comparing them against established supervised classification models (obtained through fine-tuning) from different perspectives. When faced with a severe lack of labelled samples (e.g., having only around 4 – 64 samples), the zero and few-shot prompting often outperforms fine-tuning of smaller models (Ma et al., 2023; Hongjin et al., 2022; Logan IV et al., 2022; Gao et al., 2021). However, with enough labelled samples (depending on the task, it can be as low as 100 (Pecher et al., 2024a)), the smaller models obtained through fine-tuning or instruction tuning often achieve performance on par or better than the large language models (Schick and Schütze, 2021; Qin et al., 2023). However, the behaviour in low-resource languages and other settings is not well understood, as the majority of studies focus on English and other high-resource languages.

In this work, we shift the perspective and instead of focusing on large language models as end-

task classifiers, we ask whether we **can leverage their strong generation capabilities as a form of data-driven distillation in low-resource settings**. Specifically, we investigate whether the *synthetic samples* generated by the large multilingual language models can be used to improve smaller models to achieve comparable or even superior performance in low-resource multilingual classification. Namely, we investigate three options for training or improving small models: 1) through *fine-tuning* with synthetic samples; 2) *in-context learning with smaller LLMs* using synthetic samples as in-context examples; and 3) *instruction-tuning of smaller LLMs* with synthetic samples. Afterwards, we systematically compare such small models with the large model used for generating the synthetic samples across *11 typologically diverse languages* (from low-, medium- and high-resource language groups) and *4 classification tasks*. To provide a more in-depth analysis, we also investigate how the comparison changes as the number of available synthetic samples increases and how the sensitivity to the effects of randomness affects it. The aggregated results are presented in Figure 1. Overall, within our experiments, we perform over 30,000 fine-tuning and instruction-tuning of the smaller models. Our main contributions and findings are<sup>1</sup>:

- We present a comprehensive analysis of the role of LLMs as generators of synthetic samples for low-resource classification, which can be used to train more efficient (in terms of performance and computational costs) smaller models. We perform the analysis across 11 languages and 4 classification tasks, comparing the large generation model with fine-tuning of a multilingual encoder and 3 smaller LLMs, which are used through prompting, in-context learning, and instruction-tuning.
- We show that even small amounts (20 on average for low- and medium-resource, and 100 for high-resource languages) of synthetic samples allow smaller models to match or outperform the large generator model itself, especially in low-resource settings. At the same time, performance in high-resource settings is more task-dependent, with popular and overly represented tasks often showing no performance benefit from synthetic samples, only a

---

<sup>1</sup>To support replicability, we provide the source code of our experiments at <https://github.com/kinit-sk/multilingual-classifiers-not-generators>

reduction in computation by having a smaller model. As such, it is more efficient to utilise the LLMs as generators of synthetic samples in low-resource languages that can be used to empower smaller models rather than using them as classifiers.

- The synthetic samples can help achieve similar performance to human-labelled samples in severely resource-constrained settings, but when increasing the number of available samples, their benefit quickly stagnates in comparison to the human-labelled samples, due to lower diversity and informativeness of the generated samples.

## 2 Related Work

Since their emergence, Large Language Models (LLMs), such as GPT-4 and Llama, have become central tools for generating and augmenting synthetic data. This strategy has proven effective across diverse applications, including automated scoring (Fang 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 categorization (Piedboeuf and Langlais, 2023), and personalised recommendation (Liu et al., 2024)

Although much of this progress has been centred on English, recent studies highlight the growing use of LLMs for synthetic data creation in low-resource languages. Multilingual generation has been employed to build datasets for question answering (Kramchaninova and Defauw, 2022; Namboori et al., 2023; Putri et al., 2024), fact-checking (Chung et al., 2025), commonsense reasoning in culturally nuanced settings (Pranida et al., 2025), and cross-lingual named entity recognition (Liu et al., 2021). Similarly, LLM-based data augmentation has supported sentiment stance detection (Zotova et al., 2021) and classification tasks in diverse languages (Glenn et al., 2023; Anikina et al., 2025).

Multiple studies have already focused on comparing smaller models that are further improved via fine-tuning or instruction-tuning, with larger general models, which are used through prompting or in-context learning, on text classification tasks (Dong et al., 2022; Liu et al., 2023). Many studies find that with enough labelled samples, the smaller models can outperform the significantly

larger ones (Schick and Schütze, 2021; Qin et al., 2023; Hernandez et al., 2023; Pecher et al., 2024a). The number of required samples to achieve this superior performance is often quite small, in the range between 100 – 1000 samples (Pecher et al., 2024a). Only in the extremely resource-constrained setting, where the models are fine-tuned with the same number of samples as are used for in-context learning (e.g., 4 – 64), the zero and few-shot prompting is able to outperform the smaller models (Ma et al., 2023; Hongjin et al., 2022). However, when focusing on a fairer comparison, where the models of similar sizes are used, the fine-tuning (or instruction-tuning) often leads to significantly better performance (Mosbach et al., 2023). All of the previous works perform the comparison on English or other high-resource languages. When moving to more low-resource settings, either using low-resource languages or tasks that lack labelled samples, the situation is often different. The small models can often outperform much larger LLMs even under the same data budgets, as the in-context learning often falters and is characterised by significant performance variance (Asai et al., 2024; Micallef and Borg, 2025). Similarly, when human-labelled samples are lacking, generating synthetic samples for question answering can lead to high performance, sometimes even outperforming the original training sets (Puri et al., 2020).

However, to the best of our knowledge, there are no works that investigate whether the large language models are better as generators (a task they were designed for) than classifiers, especially in a low-resource setting. In this paper, we fill this gap and ask whether we can use the large language models to generate synthetic samples in the low-resource languages that are then used to empower (e.g., using them for fine-tuning or as in-context examples) smaller models and lead to more efficient classification (in terms of performance and computation costs).

### 3 Methodology: Data-driven Distillation Through Synthetic Samples

We analyse whether it is more efficient (for classification in low-resource settings) to use the large language models as generators of synthetic samples instead of classifiers. In essence, we aim to distil the LLM into smaller models that require lower computational resources, while maintaining or even surpassing the performance of the origi-

nal LLM. To achieve this, we perform two steps: 1) generating synthetic samples; and 2) using the samples for training or as in-context examples.

**Generating Synthetic Samples** To generate the synthetic samples, we use the current state-of-the-art approaches for various languages identified in Anikina et al. (2025); Cegin et al. (2024b). The studies found that the generation of synthetic samples works significantly better when presenting the model with a few representative samples for each label and language. Therefore, we randomly sample 10 examples from the train data and include them in the prompt given to the LLM as examples from the dataset to guide the generation. Collecting such a number of human-labelled samples is also achievable in low-resource settings, while not producing as large a performance improvement as the synthetic samples that are generated using these samples. The prompt used for this generation is included in Appendix A. As part of the generated samples may be problematic, such as those of lower quality or containing different languages (mainly English), we apply self-revision as a filtering mechanism. In this step, the generator LLM is used as a judge, by presenting it with the task description, the generated synthetic sample, and the criteria (e.g., containing only text in the specified language), and asking whether the sample is acceptable. Besides relying solely on the LLM for this task, we also manually check a subset of the samples after generation, and if necessary, generate more. We generate samples until we have a total of 200 samples for each label present in the dataset.

**Using Synthetic Samples in Smaller Models** After obtaining the synthetic samples, we use them to train smaller models on the given tasks and languages. First, we use them to train a multilingual classification model through *fine-tuning*, where we initialise a pretrained model, add a classification head with dropout, and train all of the layers with a small learning rate. Second, we randomly select a set of these synthetic samples to be used as in-context examples for smaller LLMs. Third, we use synthetic samples to instruction-tune the smaller LLM as well (using LoRA) and then use it with in-context learning. The fine-tuning approach represents the cheapest option in terms of computational resources, while instruction-tuning is the most expensive. Finally, we also compare all of these approaches with a simple zero-shot prompting with small LLMs and the large LLM used for

generating the samples.

As we are also interested in analysing how the behaviour of the models changes as we use more synthetic samples, we repeat the whole process of training or improving the small models across different sizes of the generated samples. We randomly subsample 200 generated samples into subsets of different sizes, starting with 10 samples (ensuring we have at least 2 samples per class), up to the full set of samples. Furthermore, we are interested in the impact of this random subsampling, and other sources of randomness, such as a random initialisation of the models or a choice of in-context examples that were observed to cause significant variance in results in previous works (Sclar et al., 2023; Weber et al., 2023; Pecher et al., 2024b,a; Gundersen et al., 2022). As such, we run all experiments 20 times, reporting the mean accuracy and standard deviation across these runs (unless specified otherwise). More detailed setup for the models, such as prompt format, learning rate, and the sizes of the dataset subsets, is provided in Appendix A.

## 4 Experiments

**Languages and Datasets** We use 11 typologically diverse languages belonging to different groups based on their sizes. From the low-resource languages, we use Azerbaijani, Romanian, Slovenian, Telugu, and Welsh. For the medium-resource languages, we use Hebrew, Indonesian, Swahili, and Thai. Finally, for the high-resource languages, we use English and German. We evaluate all of these languages in three classification tasks of different characteristics, such as their difficulty: 1) sentiment classification; 2) topic classification; and 3) intent recognition. In addition, we evaluate English on sarcasm detection to further explore the benefit of synthetic samples not only on low-resource languages, but also on low-resource tasks in high-resource languages.

For *sentiment classification*, containing two labels, no single multilingual dataset currently exists. We therefore use the **combined ten datasets** for low-resource languages from (Gurgurov et al., 2025, 2024) with two additional datasets for English and German from (Mollanorozy et al., 2023). This combination was first introduced in (Anikina et al., 2025). These datasets vary in coverage and text domain: the German set focuses on transportation and infrastructure, while the Romanian data primarily consists of product reviews. For *topic*

*classification*, we rely on **SIB-200** (Adelani et al., 2024), which provides seven topic labels derived from the FLORES-200 machine translation corpus, annotated at the sentence level. For *intent recognition*, we utilise the **MASSIVE** dataset (FitzGerald et al., 2023), a multilingual benchmark for virtual assistant evaluation. We restrict the label set to ten distinct, most common intents to avoid possible semantic overlaps. Finally, for *sarcasm detection*, treated as binary classification in this paper, we use the **SemEval 2021 Task 7** (Meaney et al., 2021) data used for humour, offence, and sarcasm detection.

**Models** For generating the synthetic samples, we use the **LLaMA-3 70B** (Grattafiori et al., 2024) instruction-tuned large language model. We chose this model as it is open-source (and does not face problems with reproducibility, unlike many of the closed-source models) and has shown the best generation capabilities in our preliminary experiments. When comparing with other models, we use the model in a zero-shot setting. For the smaller models, we use **XLm-RoBERTa Large** (Conneau et al., 2020) for fine-tuning (due to its multilingual capabilities) and **LLaMA-3.1-8B** (Grattafiori et al., 2024), **Gemma-3-4B** (Team, 2025), and **Qwen-2.5-7B** (Team, 2024) instruction-tuned large language models for prompting, in-context learning, and instruction-tuning.

### 4.1 Comparison Across Language Groups

In this section, our goal is to answer the following main research question: **RQ1: Can we distil the large language model by generating synthetic samples and using them for training smaller multilingual models to achieve more efficient classification performance on low-resource languages?** We also specifically focus on how many labelled samples we need to generate from the large language model in order to observe similar or better performance. The results for the LLaMA-3.1-8B model, aggregated across all tasks (excluding sarcasm detection), for low-resource, medium-resource, and high-resource languages are presented in Figure 2. The results for the remaining two smaller models (Gemma-3-4B, Qwen-2.5-7B) closely follow those observed for the LLaMA-3.1-8B model and, as such, are included in Appendix B.1 (with a description of potential differences).

**Using generated synthetic samples, smaller models quickly outperform the large language**

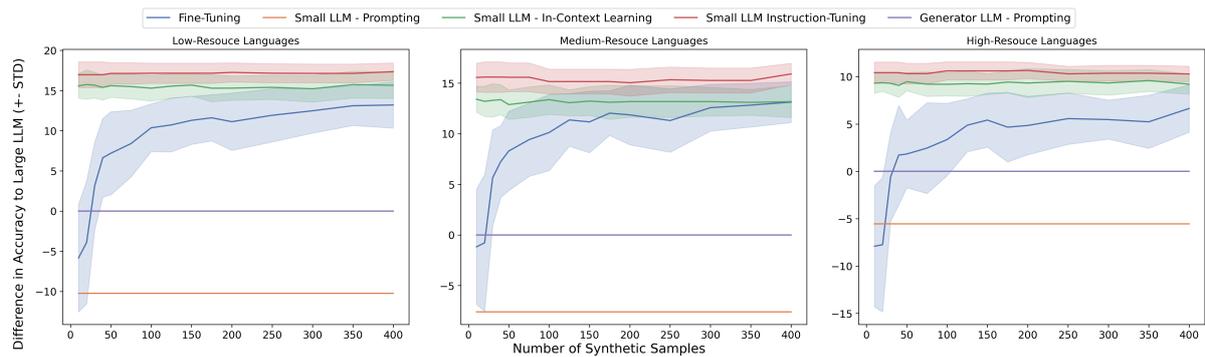


Figure 2: The difference in accuracy of smaller models trained using synthetic data compared to the large language model used for generating the data, aggregated across all tasks and language groups of different sizes. Using as few as 50 synthetic samples, the smaller models achieve higher performance across all language groups.

**model used for the generation.** To outperform the generator model, the smaller models require as few as 50 synthetic samples. When taking the randomness (represented as standard deviation over the 20 runs) originating from different subsampling of the synthetic samples and training into consideration, we observe differences only for high-resource languages, where fine-tuning requires up to 100 – 150 samples to always outperform the generator LLM. The overall increase in performance in comparison to this model, when using all the generated synthetic samples, is on average 10% (percentual points). As can be expected, the highest performance is achieved when instruction-tuning the smaller language models. However, it is significantly more expensive than fine-tuning or inference through in-context learning, while not producing as much benefit.

**The benefit of distillation through synthetic samples depends on language groups,** with low-resource languages benefiting the most and high-resource languages the least. For low-resource languages, the increases are on average 13% for fine-tuning, 15% for in-context learning, and 18% for instruction-tuning. For the high-resource languages, the increase drops to 5 – 6% for fine-tuning, 9% for in-context learning and 11% for instruction-tuning. This difference can be expected, as high-resource languages (especially English) are over-represented in all the large language models.

Overall, we can conclude that the **large language models are better suited as generators of synthetic samples that are used to train smaller models than classifiers,** especially for low-resource languages. Such a setup not only leads to better performance but also more efficient classification, as we do not require as many compu-

tational resources. However, it is important to note that we **observe a significant impact of hyperparameter setup** (represented by the large standard deviation in the results in Figure 2 due to failed runs, i.e., runs achieving performance worse than random chance) in the low-resource languages and when using synthetic samples – especially due to the tendency of the fine-tuning models to quickly overfit on the synthetic samples. This effect can be especially observed when changing the number of samples we have available, such as in the low- or high-resource languages between 150 and 200 samples, where increasing the number of available samples can lead to a drop in performance.

## 4.2 Impact of Task Characteristics

In this section, we answer the following research question **RQ2: How do the task characteristics affect the training using synthetic samples?** We analyse and compare the training across tasks with different characteristics, such as their difficulty, text lengths or how represented they are in the training data used in LLM pretraining. The results for the LLaMA-3.1-8B model are presented for the low-resource, medium-resource and high-resource languages in Figure 3. The results for the remaining models are included in Appendix B.2.

**Task characteristics have a significant effect on the success of using synthetic training samples.** For the popular sentiment classification task (which is characterised by the lower difficulty due to binary classification and over-representation in pretraining data), we observe that fine-tuning on synthetic samples cannot outperform the large model used for generating the samples. For the low-resource languages, the performance only approaches the generator model when using all of the

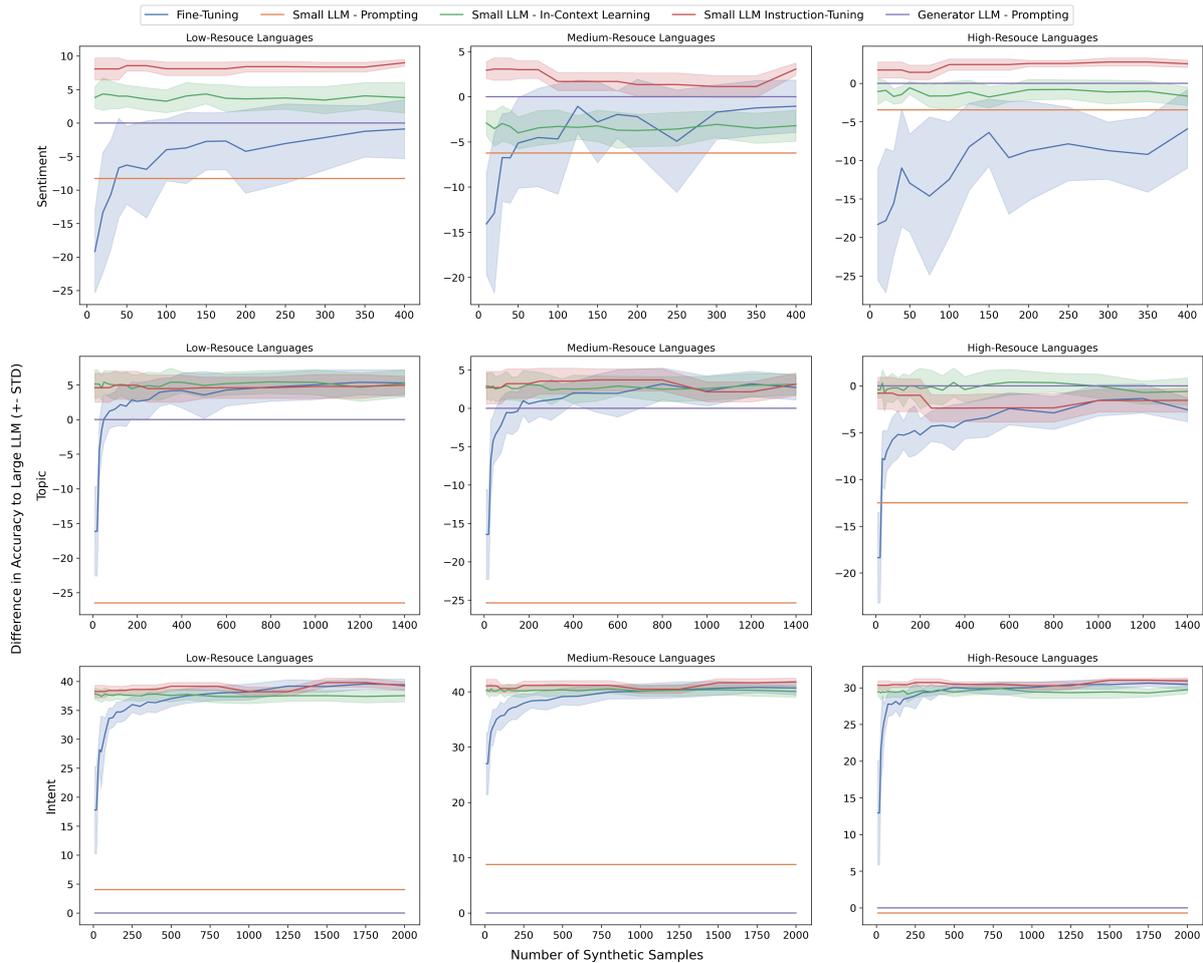


Figure 3: The comparison of the accuracy difference for tasks with different characteristics. For the commonly used sentiment classification task, we can observe significantly lower benefit of synthetic samples, especially for fine-tuning.

400 synthetic samples, while for the high-resource languages, we observe the performance of fine-tuning being lower by 5% than the generator model. This follows the findings from previous works that found that fine-tuning on overly represented tasks (sentiment classification) requires a large number of samples to outperform classification with LLMs (Pecher et al., 2024a). Similarly, in the case of medium and high-resource languages, even using synthetic in-context examples with a smaller LLM does not lead to better performance. The only consistently better approach that utilises synthetic samples is instruction-tuning. However, it also incurs significant computation costs, which can often match those from using the large generator LLM.

On the other hand, for tasks characterised by a larger number of classes, longer sentences, and lower representation in pretraining data (topic and intent classification), we observe a large benefit of synthetic samples. For the topic classification,

using synthetic samples for fine-tuning, instruction-tuning, or as in-context examples, is not able to outperform the generator LLM only in the case of high-resource languages. For medium and low-resource languages, we still observe up to 5% improvement in the performance, which is significantly (i.e., by up to 30%) higher than prompting with smaller LLMs. For the intent classification, which represents the "hardest" task (in terms of number of classes and representation), all of the smaller models utilising synthetic samples outperform the large generator model with as low as 10 samples. On low-resource languages, we observe an increase in performance of up to 40%, while on high-resource languages up to 30%.

As such, we can conclude that **generating synthetic samples and using them as distillation is more effective for less represented tasks even in low-resource languages**. This indicates that the benefit does not depend only on the resources

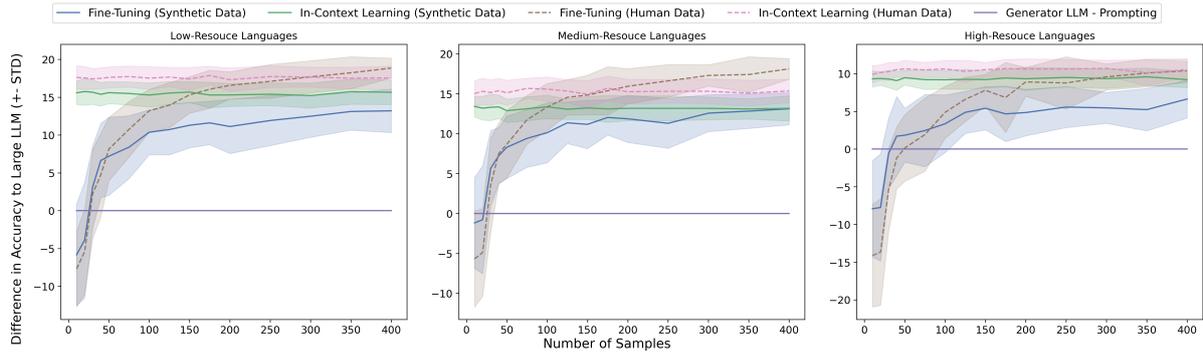


Figure 4: Comparison between synthetic and human labelled training samples. The human labelled samples provide more performance benefit, especially on larger number of samples.

available for the languages, but the representation of tasks in the pretraining data or their popularity plays a strong role as well. To further confirm this finding, we run an additional analysis on a more niche task in Section 4.4.

### 4.3 Comparison Between Synthetic and Human Data

In this section, we focus on answering the following research question **RQ3**: *How does the benefit of synthetic samples compare to human labelled data?* We use the human-labelled training samples from each language and task and use them to fine-tune the model or use them as in-context examples, following the same methodology (e.g., for subsampling, hyperparameters and overall training). We do not report results for the instruction-tuning as it incurs significant computation costs, while achieving similar performance to synthetic samples in our preliminary results. The results for the LLaMA-3.1-8B model aggregated over all tasks are presented for low-resource, medium-resource, and high-resource languages in Figure 4, and for the remaining models in Appendix B.3.

**Synthetic samples are comparable to human-labelled samples when it comes to the performance improvements only in resource-constrained settings.** When using only a small amount of samples for fine-tuning, up to 100, we observe similar performance regardless of whether using synthetic or human-labelled data. However, as the number of samples increases, the performance from human-labelled data steadily increases, while for the synthetic samples, it starts to stagnate. For low- and medium-resource languages, the gap between synthetic and human data is up to 10% when using all 400 samples, whereas for high-resource languages, the gap is only up to 5%. For

in-context learning, we observe a consistent benefit of human data over synthetic data, with a smaller gap of around 4% for low- and medium-resource languages and around 2% for high-resource languages. This follows the findings from previous works that found that human-labelled samples provide larger diversity (Guo et al., 2024; Reinhart et al., 2025) and may also explain the hyperparameter sensitivity (e.g., tendency to overfitting) when training on synthetic samples. However, it is important to note that the hyperparameter sensitivity and overfitting is observed also for the human-labelled samples (with the hyperparameters suited for synthetic samples being the best also for human-labelled ones). As the number of samples increases, the situation is improving slightly in both cases (human and synthetic samples). At the same, we do not employ any diversity incentives when generating samples, which may be the cause of lower performance benefits when using a larger number of synthetic samples (as compared to human-labelled ones) and using such incentives could help in these cases (Cegin et al., 2024a). As such, this further reinforces the finding that distillation using synthetic samples (as is) should be used mostly in resource-constrained settings (low-resource languages and niche tasks) where getting human-labelled samples is problematic. For less resource-constrained settings, more attention to sample diversity should be given to potentially observe similar benefits.

Finally, **human-labelled samples lead to lower deviation in results than synthetic samples.** For the fine-tuning with human data, we observe a significantly lower number of failed runs and reduced sensitivity to hyperparameter setup and other sources of randomness.

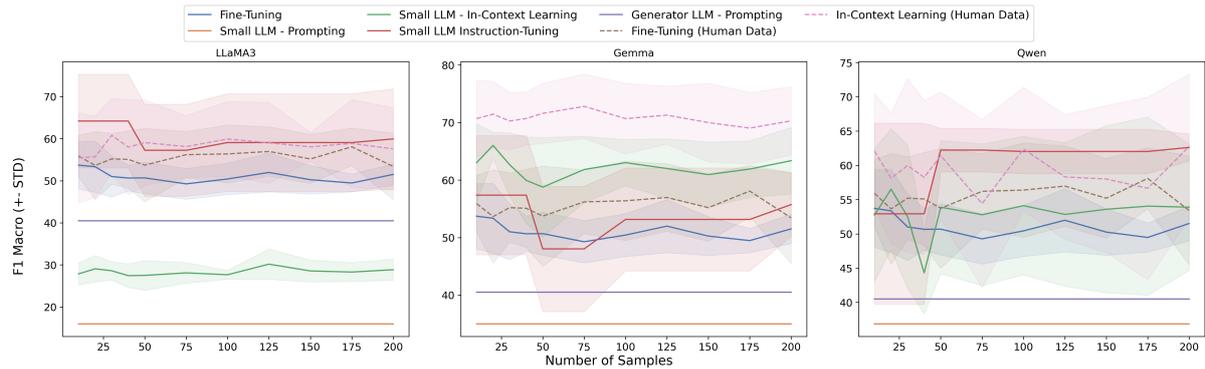


Figure 5: Using synthetic samples for sarcasm detection on English. On this more complicated task, the synthetic samples lead to performance improvements even on high-resource language.

#### 4.4 Additional Analysis: Low-Resource Task on High-Resource Language

In this section, we extend our analysis of low-resource settings to sarcasm detection in English, a more complicated task characterised by lower availability of samples for training (and lower representation in pretraining samples). Our main focus is to investigate whether the benefit of synthetic samples transfers from the low-resource languages also to a combination of a high-resource language and a more niche task. The results for all three small models are presented in Figure 5 using F1 macro due to the test set containing a significant sample imbalance (e.g., it is possible to achieve more than 80% accuracy by predicting only a single class).

**For high-resource language, using synthetic samples leads to the largest benefit when interested in more complicated or niche task.** We observe that fine-tuning can always outperform the large generator model, with a difference in F1 of 12%, even though using more samples does not lead to any noticeable increase in the overall performance. This may be caused by the inability of the generator to produce samples with sufficient informativeness. However, we observe the same behaviour for the human labelled, where the performance remains more or less stable, but higher by up to 4% over synthetic samples. For other approaches, we observe more varied results. In the case of the LLaMA-3.1-8B model, using synthetic samples in in-context learning does not lead to better performance than the generator model, although for Gemma-3-4B and Qwen-2.5-7B, it does (and leads to performance increases of 5 – 25% F1). On the other hand, using human-labelled samples leads to consistent improvements of up to 30%. Similar behaviour can be observed for the instruction-

tuning as well.

Furthermore, we notice a significantly higher variance in results across all of the approaches. It is especially noticeable for the instruction-tuning and in-context learning approaches, where the difference can be as high as 20% F1, just based on how the samples are selected. As such it further reinforces the finding that **when using synthetic samples, more focus should be dedicated to the hyperparameter setup**, such as choosing learning rate, batch size or how the samples are selected.

## 5 Conclusion

In this work, we perform a systematic and comprehensive analysis to investigate whether the large language models are better employed as classifiers or as generators of synthetic samples for training smaller models in low-resource classification. We perform a kind of data-driven distillation of a large multilingual language model, which is used to generate a set of synthetic samples that are then used to fine-tune or instruction-tune a smaller model, or used as in-context examples. Based on our experiments spanning 11 languages and 4 tasks, we show that only a small number of generated synthetic samples is enough for the smaller models to achieve performance comparable to, and often surpassing, the large generator itself, especially in low-resource languages and less represented tasks. However, we also observe potential drawbacks with this approach, where the fine-tuning on the synthetic samples shows a significantly stronger sensitivity to the hyperparameter setup. At the same time, the synthetic samples lead to improvements comparable to human-labelled data only in settings with low number of samples (around 50), with their benefit quickly stagnating due to their limited diver-

sity and informativeness. As such, our findings suggest that the large language models are best utilised as generators in low-resource settings or for the purpose of quick prototyping, producing synthetic data that effectively distils their knowledge into smaller, more efficient models for practical use, rather than classifiers. This paradigm reduces the reliance on large models at inference time and offers a scalable path toward NLP in resource-constrained settings. As a possible future extension, we see a benefit of a more sophisticated generation strategy that would provide samples with higher diversity.

## Acknowledgments

This work was partially funded by the EU NextGenerationEU through the Recovery and Resilience Plan for Slovakia under the project *AI-Auditology*, No. 09I03-03-V03-00020.

Part of the research results was obtained using the computational resources procured in the national project funded by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90254).

## Limitations

Due to the large scope of experiments in our study, focusing on multiple approaches and 11 different languages, which leads to us performing more than 30,000 fine-tuning and instruction-tuning runs, we limit the number of models used for generation to 1, fine-tuning models to 1, smaller LLMs to 3, and the number of tasks to 4. We try to mitigate this limitation by selecting the best-performing models through a preliminary study and focusing on well-established tasks.

Although we work in a setting with low-resource languages and use synthetic samples only, to generate the synthetic samples, we utilise a few human examples to drive the generation. Such an approach was found to produce significantly better examples, but may lead to samples similar to the datasets we are trying to replace. In addition, it also leads to potentially lower diversity of the samples, which we observe during our experiments. These issues could be solved by using approaches that drive the generation to create more diverse samples.

We use the same prompt for all the models, which is a result of prompt engineering on the LLaMA-3.1 model. The prompt was created based on the dataset description, the prompts used in related works, the formats recommended for differ-

ent models (e.g., taking inspiration from (Sun et al., 2023)), and our own experience. As such, this may not represent the optimal format for all the models (as identified in previous works (Sclar et al., 2023; Weber et al., 2023; Pecher et al., 2024b)), and performing the investigation on multiple different prompts may improve the overall model performance and affect the findings. However, we opted for using only a single optimised prompt format to reduce the computation costs.

## Ethical Considerations

The experiments in this paper work with publicly available datasets, citing the original authors. We do not work with any personally identifiable information or offensive content and perform no crowdsourcing for further data annotation. In addition, we are not aware of any potential ethical harms or negative societal impacts of our work, apart from the ones related to the advancement of the field of machine learning. We follow the license terms for all the models we use (such as the one required for the use of the LLaMA-3 models). It is possible that the large language models we use contain biases and potentially offensive or harmful content. However, the original authors of these models reduce this bias as much as possible. At the same time, although we generate synthetic samples and include them as part of the supplementary material, we believe that no problematic samples are contained due to the tasks we have chosen. Finally, as the highest impact of our study is the CO<sub>2</sub> generated as part of our experiments, we publicly release the impact statement and the number of computers used below.

**Impact Statement: CO<sub>2</sub> Emissions Related to Experiments** The experiments presented in this paper used significant computational resources as they required multiple training and evaluation runs of multiple models (to deal with variance in results), as well as using large language models that require a lot of computation even just for the inference. Overall, the experiments were conducted using a private infrastructure, which has a carbon efficiency of 0.432 kgCO<sub>2</sub>eq/kWh (default value used as the actual efficiency of our HW instance was not measured). A cumulative of 1000 hours of computation was performed on hardware of type A100 PCIe 80GB (TDP of 250W). Total emissions are estimated to be 108 kgCO<sub>2</sub>eq of which 0 percents were directly offset. This includes all of the gener-

ation, training, and inference with all models, and the preliminary experiments. These estimations were conducted using the [Machine Learning Impact calculator](#) presented in (Lacoste et al., 2019).

Whenever possible, we tried to reduce the compute resources used as much as possible. Most computational resources were used by the instruction-tuning of the smaller LLMs. As such, we decided to perform this operation with LoRA, while also limiting the number of steps, to reduce the cost as much as possible. Similarly, we do not perform the experiments across various large language models for generation, for fine-tuning, and as small LLMs, but instead choose only the best performing ones, to limit the number of training and evaluation runs we need in this paper, while also providing a good enough scope.

## References

- 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.
- Tatiana Anikina, Jan Cegin, Jakub Simko, and Simon Ostermann. 2025. [A rigorous evaluation of llm data generation strategies for low-resource languages](#). *Preprint*, arXiv:2506.12158.
- Akari Asai, Sneha Kudugunta, Xinyan Yu, Terra Blevins, Hila Gonen, Machel Reid, Yulia Tsvetkov, Sebastian Ruder, and Hannaneh Hajishirzi. 2024. [BUFFET: Benchmarking large language models for few-shot cross-lingual transfer](#). 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 1771–1800, Mexico City, Mexico. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901.
- 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.
- 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. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. [A survey for in-context learning](#). *arXiv preprint arXiv:2301.00234*.
- Luyang Fang, Gyeong-Geon Lee, and Xiaoming Zhai. 2023. [Using gpt-4 to augment unbalanced data for automatic scoring](#). *Preprint*, arXiv:2310.18365.
- 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 Nataraajan. 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.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). 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 3816–3830, Online. ACL.

- 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. -.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Odd Erik Gundersen, Kevin Coakley, and Christine Kirkpatrick. 2022. Sources of irreproducibility in machine learning: A review. *arXiv preprint arXiv:2204.07610*.
- Yanzhu Guo, Guokan Shang, and Chloé Clavel. 2024. Benchmarking linguistic diversity of large language models. *arXiv preprint arXiv:2412.10271*.
- Daniil Gurgurov, Mareike Hartmann, and Simon Oostermann. 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 Oostermann. 2025. [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.
- Evan Hernandez, Diwakar Mahajan, Jonas Wulff, Micah J Smith, Zachary Ziegler, Daniel Nadler, Peter Szolovits, Alistair Johnson, Emily Alsentzer, and 1 others. 2023. Do we still need clinical language models? In *Conference on Health, Inference, and Learning*, pages 578–597. PMLR.
- SU Hongjin, Jungo Kasai, Chen Henry Wu, Weijia Shi, Tianlu Wang, Jiayi Xin, Rui Zhang, Mari Ostendorf, Luke Zettlemoyer, Noah A Smith, and 1 others. 2022. Selective annotation makes language models better few-shot learners. In *The Eleventh International Conference on Learning Representations*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- 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.
- Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700*.
- 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.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ACM Comput. Surv.*, 55(9).
- 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.
- Robert Logan IV, Ivana Balazevic, Eric Wallace, Fabio Petroni, Sameer Singh, and Sebastian Riedel. 2022. [Cutting down on prompts and parameters: Simple few-shot learning with language models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2824–2835, Dublin, Ireland. ACL.
- Yubo Ma, Yixin Cao, Yong Hong, and Aixin Sun. 2023. [Large language model is not a good few-shot information extractor, but a good reranker for hard samples!](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10572–10601, Singapore. ACL.
- J. A. Meaney, Steven Wilson, Luis Chiruzzo, Adam Lopez, and Walid Magdy. 2021. [SemEval 2021 task 7: HaHackathon, detecting and rating humor and offense](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 105–119, Online. Association for Computational Linguistics.
- Kurt Micallef and Claudia Borg. 2025. [Melabenchv1: Benchmarking large language models against smaller fine-tuned models for low-resource maltese nlp](#). *arXiv preprint arXiv:2506.04385*.
- 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.
- Marius Mosbach, Tiago Pimentel, Shauli Ravfogel, Dietrich Klakow, and Yanai Elazar. 2023. [Few-shot](#)

- fine-tuning vs. in-context learning: A fair comparison and evaluation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 12284–12314, Toronto, Canada. ACL.
- 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.
- Branislav Pecher, Ivan Srba, and Maria Bielikova. 2024a. Comparing specialised small and general large language models on text classification: 100 labelled samples to achieve break-even performance. *arXiv preprint arXiv:2402.12819*.
- Branislav Pecher, Ivan Srba, and Maria Bielikova. 2024b. A survey on stability of learning with limited labelled data and its sensitivity to the effects of randomness. *ACM Computing Surveys*, 57(1).
- 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 Zahirah Pranida, Rifo Ahmad Genadi, and Fajri Koto. 2025. Synthetic data generation for culturally nuanced commonsense reasoning in low-resource languages. *Preprint*, arXiv:2502.12932.
- Raul Puri, Ryan Spring, Mohammad Shoeybi, Mostofa Patwary, and Bryan Catanzaro. 2020. Training question answering models from synthetic data. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5811–5826, Online. Association for Computational Linguistics.
- Rifki Afina Putri, Faiz Ghifari Haznitrana, Dea Adhista, and Alice Oh. 2024. Can LLM generate culturally relevant commonsense QA data? case study in Indonesian and Sundanese. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 20571–20590, Miami, Florida, USA. Association for Computational Linguistics.
- Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao Chen, Michihiro Yasunaga, and Diyi Yang. 2023. Is ChatGPT a general-purpose natural language processing task solver? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1339–1384, Singapore. ACL.
- Alex Reinhart, Ben Markey, Michael Laudenschlager, Kachatur Pantusen, Ronald Yurko, Gordon Weinberg, and David West Brown. 2025. Do llms write like humans? variation in grammatical and rhetorical styles. *Proceedings of the National Academy of Sciences*, 122(8):e242245122.
- 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.
- Timo Schick and Hinrich Schütze. 2021. It’s not just size that matters: Small language models are also few-shot learners. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352, Online. ACL.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2023. Quantifying Language Models’ Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting.
- 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.
- Xiaofei Sun, Linfeng Dong, Xiaoya Li, Zhen Wan, Shuhe Wang, Tianwei Zhang, Jiwei Li, Fei Cheng, Lingjuan Lyu, Fei Wu, and 1 others. 2023. Pushing the limits of chatgpt on nlp tasks. *arXiv preprint arXiv:2306.09719*.
- Gemma Team. 2025. Gemma 3.
- Qwen Team. 2024. Qwen2.5: A party of foundation models.
- Lucas Weber, Elia Bruni, and Dieuwke Hupkes. 2023. Mind the instructions: a holistic evaluation of consistency and interactions in prompt-based learning. In *Proc. of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 294–313, Singapore. Association for Computational Linguistics.
- 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 Experimental Setup: Further Details

In this Appendix, we provide further details on our experimental setup.

**Languages and Datasets** We use 11 typologically diverse languages belonging to different groups based on their sizes: 1) low-resource, including Azerbaijani, Romanian, Slovenian, Telugu, and Welsh; 2) medium-resource, including Hebrew, Indonesian, Swahili, and Thai; and 3) high-resource, including English and German. For each language and task, we generate 200 samples per class, which results in 400 samples for sentiment classification, 1400 for topic classification, 2000 for intent classification, and 200 for sarcasm detection (where we generate only 100 samples per class). When subsampling the datasets, we start with overall 10 samples (ensuring we have at least 2 samples per class) and increasing the number of available samples until the full dataset – the exact subsamples are 10, 20, 30, 40, 50, 75, 100, 125, 150, 175, 200, 250, 300, 350, 400, 500, 600, 800, 1000, 1250, 1500, 1750, 2000. As we perform the subsampling by randomly choosing the same number of samples from all classes, we repeat all of the experiments 20 times to deal with any effects of randomness.

**Models and Training** For fine-tuning, we use the XLM-RoBERTa large model from huggingface<sup>2</sup>. For the small LLMs, we use LLaMA-3.1 8B<sup>3</sup>, Gemma-3 4B<sup>4</sup> and Qwen-2.5 8B<sup>5</sup> all instruction tuned versions. For the large generator model, we run only zero-shot classification through prompting, while for the smaller LLMs, we also include in-context learning. The prompts used for all LLMs are included in Table 1. For in-context learning, we also evaluate different numbers of in-context examples per class, starting with 2 and increasing them up to 100, always reporting only the best performance. However, we observe that in almost all cases, the best performance is achieved using 2 samples per class (and for some, using 4). For instruction-tuning, we use the same prompt format and utilise LoRa (Hu et al., 2022), which is a method of parameter-efficient fine-tuning. We

<sup>2</sup><https://huggingface.co/FacebookAI/xlm-roberta-large>

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

<sup>4</sup><https://huggingface.co/google/gemma-3-4b-it>

<sup>5</sup><https://huggingface.co/unsloth/Qwen2.5-7B-Instruct>

set the LoRA rank to 16, alpha to 16, and dropout to 0.05. We reparametrize each key, value, and downsample weight matrix of the transformer architecture. We train the model for 600 steps with a learning rate of 3e-5, with 10 warmup steps and batch size of 4 with early stopping. After training the model, we perform inference using in-context learning. For all of prompting and in-context learning inference, we set the LLMs to produce deterministic output (by setting temperature and turning off sampling) with up to 20 samples. For fine-tuning, we observe a significant sensitivity to the hyperparameters, especially across different dataset subsamples. As such, we run a hyperparameter optimisation for each language and dataset size and, to guarantee fairness in comparison, choose the setup that works the best across all of the experimental setups. We use a dropout of 0.3 in the classification layer, a learning rate of 1e-5, an AdamW optimiser with warmup, and train for 50 epochs with early stopping. We dynamically change the batch size across different dataset sizes, starting with a batch size of 4 at 10-20 samples, 8 at 20-50, 12 at 50-100, 16 at 100-150, 32 at 150-250, and 64 for all other sizes. Without this setup, we have observed a severe overfitting of the fine-tuning models.

**Generation** For generating the synthetic samples, we use Llama-3-70b-instruct<sup>6</sup>. The generation prompt is included in Table 1. The *task\_text\_type* placeholder had values based on tasks, either semantic analysis, intent recognition, topic classification, or sarcasm detection. The *sep\_token* was represented as "—". The *label* placeholder was replaced for each task and label with an LLM-generated explanation of that label based on randomly preselected human examples. The *examples* placeholder contained the human examples for the given label being generated. We use the following parameters for generation *temperature=0.7, top\_p=0.9, max\_tokens=4096, repetition\_penalty=1.2*. We collected 10 generated samples per inference run for increased efficiency and performed cleaning of the generated data. We excluded duplicates and collected until 200 unique texts were generated.

<sup>6</sup><https://huggingface.co/TechXGenus/Meta-Llama-3-70B-Instruct-GPTQ>

Table 1: Prompt format for the large language model used for generating the synthetic samples and for the classification using prompting or in-context learning. The parts included in "{}" are replaced by the relevant data, such as task values or examples..

Task Type	Prompt Format Text
Generation	Please create 6 different $\{task\_text\_type\}$ in the $\{language\}$ language, separated by the $sep\_token$ token. The $\{task\_text\_type\}$ should be about the $\{label\}$ . Note that some examples from the dataset look as follows: Examples: $\{examples\}$ . Output only the text in $\{language\}$ and nothing else! Do not number the texts!
Sentiment	Determine the sentiment of the sentence using the following options: 1) negative; 2) positive. Use only these two options.
Topic	Determine the topic of the sentence using the following options: 1) science and technology; 2) travel; 3) politics; 4) sports; 5) health; 6) entertainment; 7) geography. Use only these options.
Intent	Determine topic of the sentence using the following options: 1) alarm query; 2) audio volume down; 3) calendar remove; 4) cooking recipe; 5) datetime convert; 6) send email; 7) play audiobook; 8) movie recommendation; 9) transport ticket; 10) weather query. Use only these options.
Sarcasm	Determine whether the comment in the sentence is sarcastic using the following options: 1) serious; 2) sarcastic. Use only these two options.

## B Results for Remaining Models

In this Appendix, we present the results of the experiments from the main part of the paper for the remaining models, specifically Gemma-3-4B and Qwen-2.5-7B, and draw a comparison with the results from the main content of the paper.

### B.1 Comparison Across Language Groups

In this section, we present the results from the main comparison, aggregated across all tasks and languages, for the Gemma (Figure 6) and Qwen (Figure 7) models. Although the results follow the ones from the main part of the paper, there are some slight differences. The smaller LLMs still benefit from using synthetic samples, but we observe smaller performance improvements as compared to the LLaMA3 model. For the Gemma model, the improvement is up to 12% for in-context learning and 10% for instruction tuning. For the Qwen model, the improvement is up to 7% for in-context learning and 10% for instruction tuning. The difference may be explained by the lower overall performance capabilities of these models, although when looking at the results from zero-shot prompting, these models often perform better than LLaMA3. As such, these models may benefit less from the synthetic samples due to the discrepancy between the generator model and smaller LLMs – the samples generated by the large LLaMA model may represent a different distribution for the Gemma and Qwen models, and thus provide lower benefit.

### B.2 Impact of Task Characteristics

In this section, we present the results for the impact of task characteristics, aggregated for the different

language groups, for the Gemma (Figure 8) and Qwen (Figure 9) models. In comparison to the results from the main content of the paper, we observe a lower benefit of the synthetic samples for the smaller LLMs used through in-context learning or instruction-tuning. The smaller models fail to outperform the generator model only in 4 cases out of 18, with both models on the sentiment task for medium-resource languages and the topic task or high-resource languages. For intent task, we do not observe such behaviour. At the same time, they often underperform the fine-tuning with synthetic samples. This follows the findings from the previous section and, as such, may be explained by the worse overall performance of these models (caused by lower number of parameters) or the discrepancy between the generator and these classification models (the samples generated by the LLaMA generator model may be better suited for the smaller LLaMA architectures as compared to different model families). However, it is important to note that the Gemma model benefits more from an increased number of samples when it comes to instruction tuning, which we do not observe for the other models. As such, it may indicate that this model has a larger room for improvement, as it is also the smallest LLM we use. Especially for the high-resource languages, the model can be improved by up to 10 – 15% through instruction tuning on the synthetic samples.

### B.3 Comparison Between Synthetic and Human Data

In this section, we present the results from comparing synthetic and human data for the Gemma (Figure 10) and Qwen (Figure 11) models. As op-

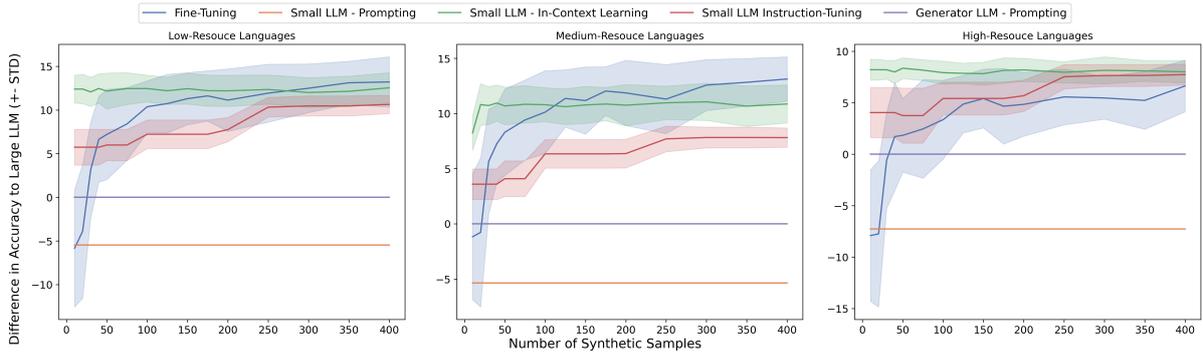


Figure 6: The difference in accuracy of smaller models (for Gemma model) trained using synthetic data compared to the large language model used for generating the data, aggregated across all tasks and language groups of different sizes.

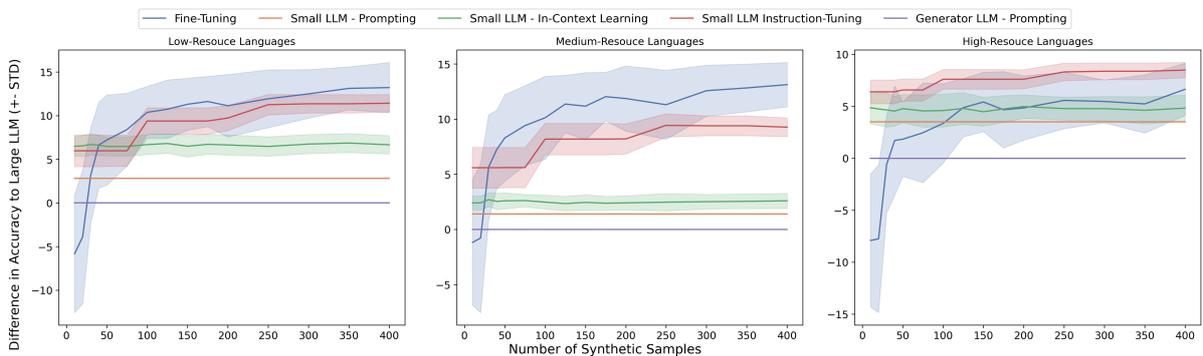


Figure 7: The difference in accuracy of smaller models (for Qwen model) trained using synthetic data compared to the large language model used for generating the data, aggregated across all tasks and language groups of different sizes.

posed to the results from the main content of the paper, we observe a significantly lower difference between synthetic and human data, with the difference as high as 2% for both the Gemma and Qwen models. This observation may also provide an explanation for the results in the previous sections, as the models benefit less from additional samples and may be less suited for in-context learning and further instruction tuning. As with the previous sections, the models are often still outperformed by simple fine-tuning.

### C Computational Costs

In this appendix, we provide a summary of the running time (determined as GPU hours) for a single run, including inference and training if necessary, on the same GPU for the models. The summary is provided in Table 2. As the number of samples that is considered for inference increases, the cost for the generator model increases at a significantly higher rate compared to other models. At the same time, the cost for instruction-tuning is mostly for

the training. We can see that using smaller models is more efficient in the long run when it comes to GPU hours required.

Table 2: Comparison of costs, calculated as the average GPU wall-clock time for a single full run, including inference and training if necessary.

MODEL	GPU TIME (H)
Generator Model	15
XML-RoBERTa	0.167
Prompting	1-2
In-Context Learning	3-4
Instruction-Tuning	10

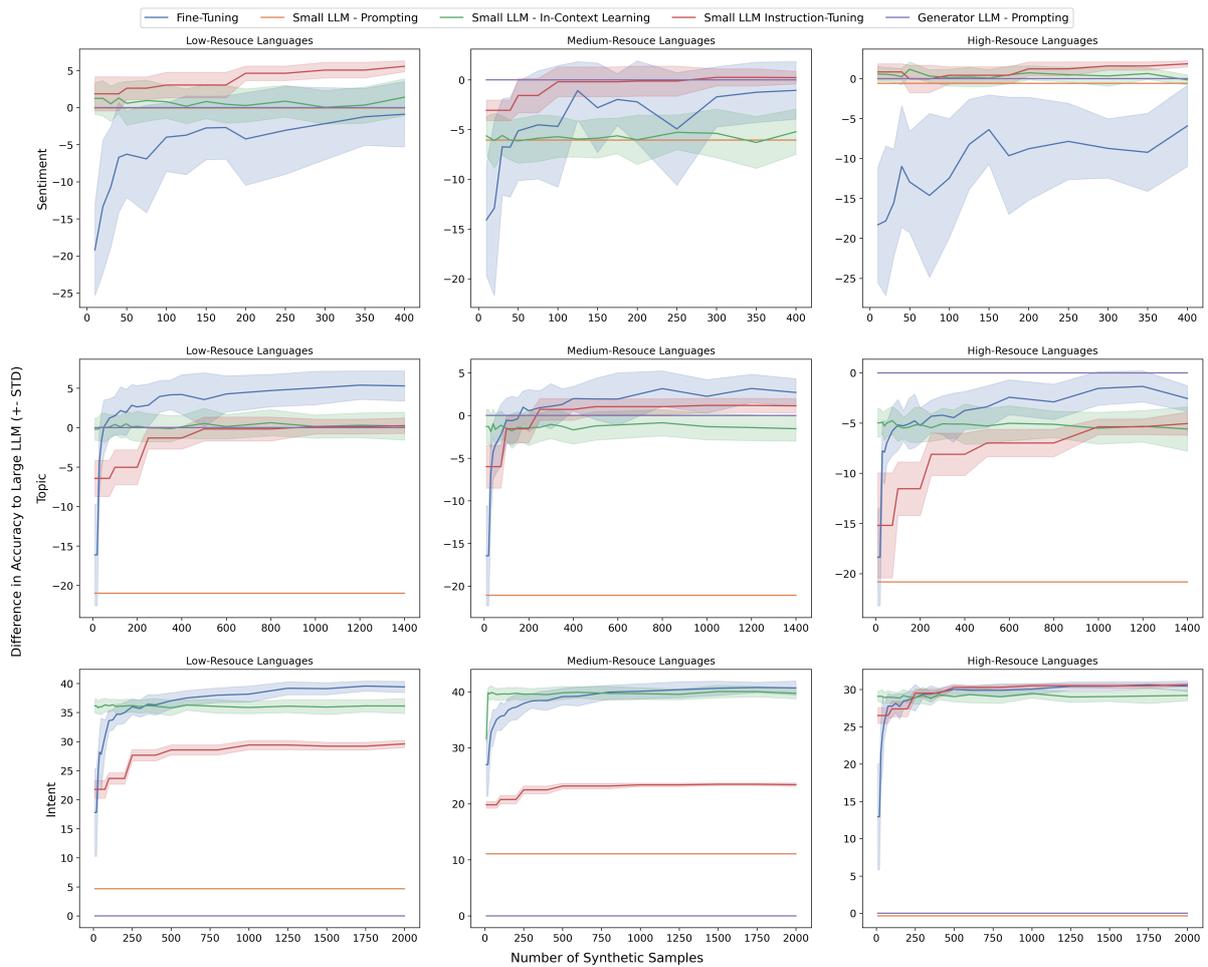


Figure 8: The comparison of the accuracy difference for tasks with different characteristics for the Gemma model.

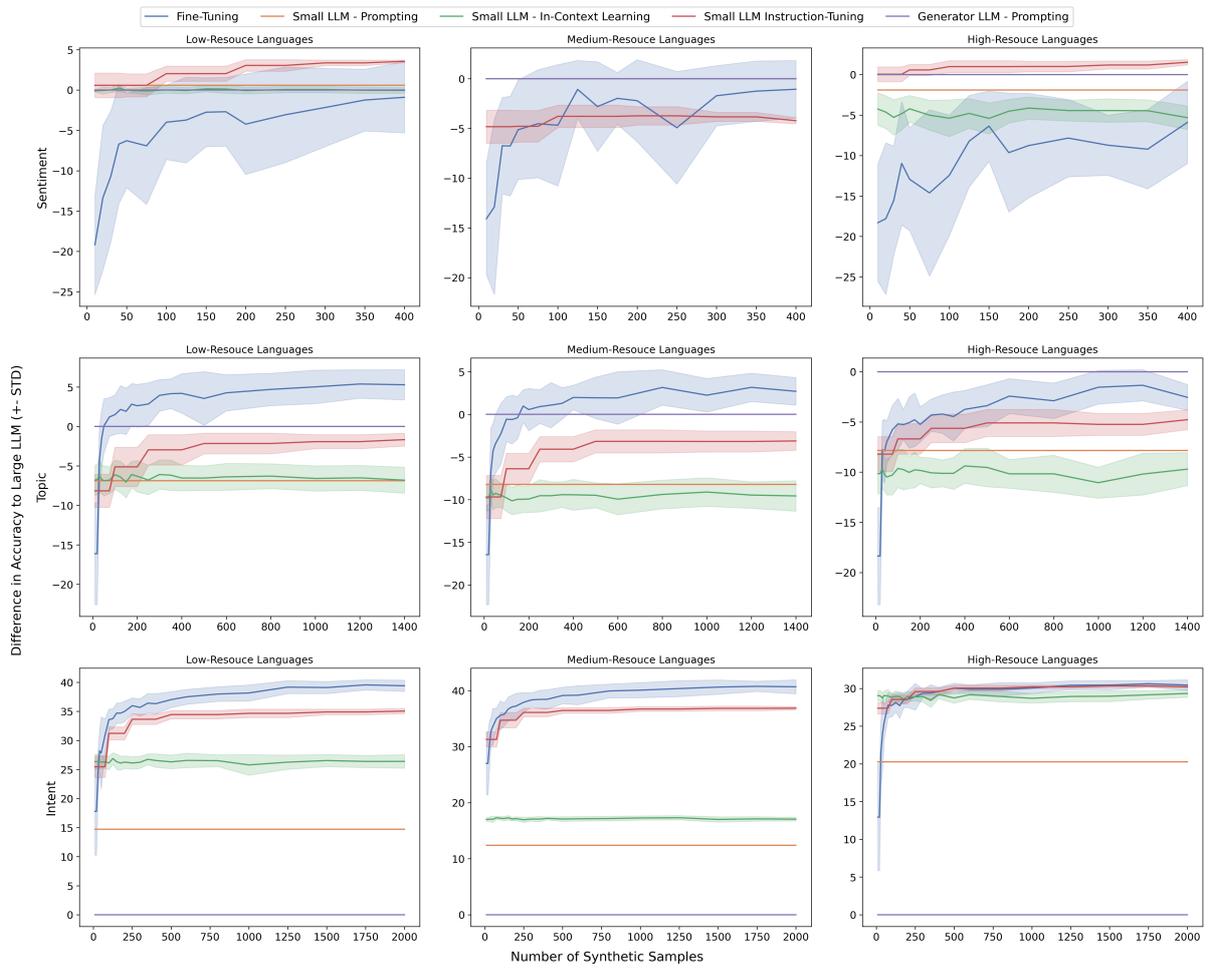


Figure 9: The comparison of the accuracy difference for tasks with different characteristics for the Qwen model.

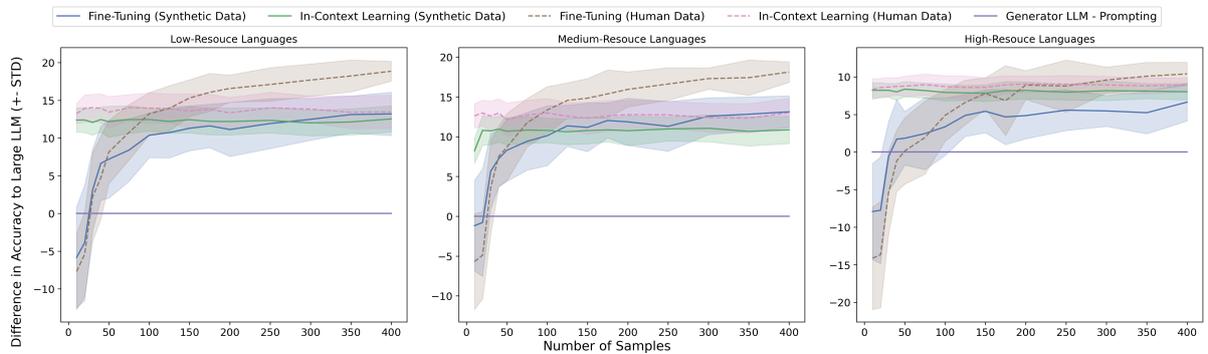


Figure 10: Comparison between synthetic and human labelled training samples for the Gemma model.

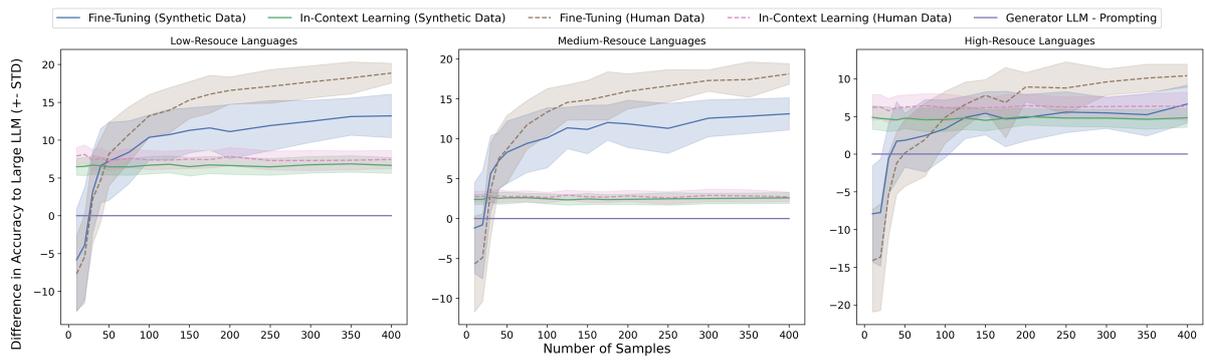


Figure 11: Comparison between synthetic and human labelled training samples for the Qwen model.