

# Split and Rephrase with Large Language Models

David Ponce<sup>\*1,2</sup> and Thierry Etchegoyhen<sup>\*1</sup> and Jesús Calleja<sup>1,2</sup> and Harritxu Gete<sup>1,2</sup>

<sup>1</sup> Fundación Vicomtech, Basque Research and Technology Alliance (BRTA)

<sup>2</sup> University of the Basque Country UPV/EHU

{adponce, tetchegoyhen, jcalleja, hgete}@vicomtech.org

## Abstract

The Split and Rephrase (SPRP) task, which consists in splitting complex sentences into a sequence of shorter grammatical sentences, while preserving the original meaning, can facilitate the processing of complex texts for humans and machines alike. It is also a valuable testbed to evaluate natural language processing models, as it requires modelling complex grammatical aspects. In this work, we evaluate large language models on the task, showing that they can provide large improvements over the state of the art on the main metrics, although still lagging in terms of splitting compliance. Results from two human evaluations further support the conclusions drawn from automated metric results. We provide a comprehensive study that includes prompting variants, domain shift, fine-tuned pretrained language models of varying parameter size and training data volumes, contrasted with both zero-shot and few-shot approaches on instruction-tuned language models. Although the latter were markedly outperformed by fine-tuned models, they may constitute a reasonable off-the-shelf alternative. Our results provide a fine-grained analysis of the potential and limitations of large language models for SPRP, with significant improvements achievable using relatively small amounts of training data and model parameters overall, and remaining limitations for all models on the task.

## 1 Introduction

Transforming complex sentences into a sequence of shorter sentences, while preserving the original meaning, can facilitate the processing of complex texts for humans and machines alike. It is, for instance, an important operation in text simplification (Zhu et al., 2010; Narayan and Gardent, 2014; Shardlow, 2014), where simpler split sentences can provide content that is easier to read and comprehend. This task, commonly referred to as Split

and Rephrase (SPRP) (Narayan et al., 2017), also provides a relevant framework to evaluate natural language processing models, as complex linguistic properties need to be tackled to provide optimal splits and appropriate rephrasing.

The SPRP task is typically performed via dedicated transformation models, e.g. sequence-to-sequence modelling with a copy mechanism (Aharoni and Goldberg, 2018) or graph-based neural segmentation (Gao et al., 2021). The main limitation for this type of approach is the dependency on training datasets of paired complex and split sentences, which are scarce across languages and domains. As an alternative, Kim et al. (2021) exploit one-to-many alignments in parallel corpora and pivot machine translation (MT) to create an aligned corpus of complex and split sentences, which can in turn be exploited by sequence-to-sequence models. For this type of approach, the main bottleneck is the dependency on parallel corpora with sufficient many-to-one alignments and on quality pivot machine translation models.

Large language models (LLMs), such as the GPT models (Radford et al., 2018) based on the Transformer architecture (Vaswani et al., 2017), have demonstrated their strong potential on a large number of downstream tasks (Radford et al., 2019; Brown et al., 2020). In this work, we measure their performance on the SPRP task, prompting the models to generate split and rephrase hypotheses from complex sentences. We evaluate different variants of the approach on datasets of varying complexity, contrasting zero shot, few shot in-context learning, and fine-tuning scenarios. We notably measure the impact of parameter and training data size, domain shift, and prompt optimisation.

Overall, our results demonstrate that LLMs outperform the previous state of the art by large margins across multiple metrics, with remaining relative deficiencies, notably in terms of split generation. Results from two human evaluations, a

\*Equal contribution.

comparative 3-way ranking task and a qualitative evaluation, further support our conclusions. Our main contributions can be summarised as follows:

- New state of the art on the SPRP task, established on publicly available datasets of varying complexity, over multiple metrics and two human evaluations.
- A comprehensive evaluation of LLMs for SPRP, covering LLM variants with and without instruction tuning, along with the impact of different prompting strategies, parameter size, training data volumes, and domain shift.
- Empirical results on LLM strengths and limitations for SPRP, in particular over splitting compliance and rephrasing variability.

## 2 Related Work

The Split and Rephrase task was first proposed by [Narayan et al. \(2017\)](#), who created the WebSplit dataset from RDF semantic tuples and provided benchmarks with statistical MT models and LSTMs with attention. [Aharoni and Goldberg \(2018\)](#) identified data representation in this dataset and provided enhancements along with a new benchmark augmenting the LSTM with a copy mechanism ([Gu et al., 2016](#)). To address some limitations of WebSplit, notably in terms of vocabulary and naturality, [Botha et al. \(2018\)](#) presented WikiSplit, a corpus created using Wikipedia data that takes advantages of edits made on a page to split or fuse sentences. They showed that models trained on WikiSplit improved those trained on WebSplit in terms of BLEU ([Papineni et al., 2002](#)). However, they also noted that WikiSplit is a relatively noisy dataset, not suited for evaluation.

Among other approaches, [Das et al. \(2018\)](#) leveraged dependency parsing information for the SPRP task, whereas [Niklaus et al. \(2019\)](#) employed hand-crafted rules, recursively applied to complex sentences to generate simpler ones. Recent approaches have exploited neural modelling, in particular the Transformer architecture. Thus, [Malmi et al. \(2019\)](#) proposed a tagging approach using a BERT model ([Devlin et al., 2019](#)) as encoder and a single Transformer decoder layer. [Stahlberg and Kumar \(2020\)](#) also explored a tagging model, at the span level instead of the token level, achieving improvements with Pointer Networks ([Vinyals et al., 2015](#)) for end of span prediction and a two-layer Transformer decoder. [Berlanga Neto and Ruiz \(2021\)](#) introduced

a cross-lingual model based on part-of-speech tags, GRU sequence-to-sequence modelling, and BERT for next word prediction. Recently, [AlAjlouni and Li \(2023\)](#) reported improvements via transfer knowledge of data generated by the DISSIM rule-based system ([Niklaus et al., 2019](#)) to fine-tune an encoder-decoder T5 model.

In this work, we compare our LLM-based variants with two recent approaches and associated corpora. First, we address BiSECT ([Kim et al., 2021](#)), a method and corpus based on extracting 1-2 alignments from parallel bilingual corpora and pivot machine translation. BiSECT achieved higher quality than WikiSplit, and sequence-to-sequence models trained on the corpus improved on the SPRP task. Additionally, we compare our work to [Gao et al. \(2021\)](#), who proposed the ABCD framework where SPRP is performed via neural graph segmentation, according to the number of predicates in the complex sentence, approaching or outperforming parsing-based and encoder-decoder models. They also provided two new datasets for the SPRP task: MinWiki, a reduced version of WikiSplit where only pairs of complex and simple sentences with a matching number of predicates are kept, and DeSSE, a corpus of students’ opinion essays in social science consisting of complex sentences and corresponding splits; we used the latter in our experiments, in addition to BiSECT.

Finally, we used the Wiki-BM corpus described in [Zhang et al. \(2020a\)](#), who found that the standard WebSplit dataset contained easily identifiable syntactic cues, which even a simple rule-based approach could exploit to match the performance of state-of-the-art models. Their benchmarks feature higher syntactic diversity to overcome this issue and are thus suited for our qualitative evaluation of LLM-based SPRP hypotheses.

## 3 Methodology

The overall approach to SPRP with LLM is simple: probe a LLM to generate sequences of meaning-preserving sentences from a single complex sentence. We explored two main variants for the task.

We first fine-tuned pretrained LLMs on pairs of complex and split sentences. The objective in this scenario is twofold: (i) measure the performance of bare LLMs on the SPRP task, and (ii) evaluate LLM accuracy with varying number of parameters. For these models, we also evaluate different approaches to prompting, comparing soft and man-

Corpus	Version	Train	Dev	Test
DeSSE	Original	13,199	42	790
DeSSE	Selected	6,432	2,000	504
BiSECT	Original	928,440	9,079	583
BiSECT	Selected	6,432	2,000	580

Table 1: DeSSE and BiSECT English corpora statistics

ually crafted prompts.

Additionally, we evaluate scenarios where we directly prompt LLMs that have been previously tuned on instruction datasets. The goal in this case is to measure the ability of this type of readily-trained models to perform the two tasks at hand, with or without positive examples. Under this scenario, we explore both zero-shot and in-context few-shot task completion, by directly prompting the model for completion in the first case, and providing a limited number of examples within the prompt in the latter case.

## 4 Experimental Setup

### 4.1 Corpora

To train and evaluate our models we first selected the aforementioned English datasets: DeSSE and BiSECT.<sup>1</sup> Statistics for each corpus are provided in Table 1, in terms of number of pairs of complex and split sentences. We indicate the figures for the original dataset (*Original*) and the versions we employed in this work (*Selected*).

We excluded all pairs where the split sentence did not contain at least one inner end-of-sentence punctuation mark (286 instances in DeSSE and 3 in BiSECT). For training, we down-sampled BiSECT to the size of DeSSE for comparison purposes. As development data, since the original sets for DeSSE were rather small (42 examples), we sampled 2,000 pairs from both training datasets.

We chose these two corpora as they present different challenges for SPRP. In terms of topics, DeSSE is based on social science opinion essays and is therefore less varied than BiSECT, which was created from a large variety of parallel corpora. In terms of complex sentence average length, the two datasets are also markedly different: for training, the average is 40.14 tokens in BiSECT,

<sup>1</sup>We use the datasets made available in the following repositories, as of April 2023: <https://github.com/serenayj/DeSSE> and <https://github.com/mounicam/BiSECT>.

contrasting with 25.25 in DeSSE; in the test sets, the averages were 39.55 and 24.10, respectively.

Finally, for our qualitative manual evaluation, we used the Wiki-BM corpus [Zhang et al. \(2020a\)](#), as it was manually crafted for the SPRP task.

### 4.2 Models

**Pretrained language models.** As pretrained LLMs, we used the Pythia suite ([Biderman et al., 2023](#)), a collection of Transformer-based decoder-only autoregressive generative language models available in different parameter sizes and trained on the Pile dataset ([Gao et al., 2020](#)). We selected the following variants for our experiments, in terms of number of model parameters: 410M, 1B, 1.4B, 2.8B, 6.9B and 12B.

**Instruction-tuned language models.** As LLMs readily tuned for instruction-based prompting, we selected Dolly 2.0,<sup>2</sup> with 7B parameters, which is based on the 6.9B Pythia language model and was tuned on 15K instructions.<sup>3</sup> This model allowed for a direct comparison between the fine-tuned LLM and the instruction-based derived model. Additionally, we selected the Alpaca-Lora model with 7B parameters, based on LLaMA ([Touvron et al., 2023](#)) and fine-tuned on the Alpaca instruction dataset,<sup>4</sup> representative of a model derived from a different LLM and set of instructions.

**Prompting.** We used different prompting strategies, depending on the model variant. For bare LLMs, we first compare the use of a simple hand-crafted prompt, namely:

INSTRUCTION: *Split the sentence below into separate sentences.*

INPUT:  $\alpha$ .

RESPONSE:  $\beta$

, where  $\alpha$  and  $\beta$  are provided at training time and the model is tasked to complete  $\beta$  at inference time. For this prompting approach, we fine-tuned the Pythia models with LoRA ([Hu et al., 2022](#)), using default parameters (see Appendix A).

Additionally, we experimented with soft-prompting methods, which have been shown to improve LLM performance on several tasks, addressing the limitations of manual prompt design.

<sup>2</sup><https://huggingface.co/databricks/dolly-v2-7b>

<sup>3</sup><https://huggingface.co/datasets/databricks/databricks-dolly-15k>

<sup>4</sup><https://github.com/gururise/AlpacaDataCleaned>

We selected two such approaches, namely prompt-tuning (Lester et al., 2021) and P-tuning (Liu et al., 2023). In both cases, trainable prompt embeddings are added to a frozen base model, with P-tuning also involving a dedicated encoder.

For zero-shot probing, we used the hand-crafted prompt described above, leaving the response empty for the model to complete. For in-context learning, we used a few-shot approach based on  $n$  examples prepended to the prompt. Thus, each query is composed of  $n$  concatenated prompts that follow the previously described pattern, along with the last prompt where the response is left empty. For our experiments, we selected the first 5 examples from the development set as examples, although other variants are of course possible.

Preliminary experiments with alternative zero-shot prompt variants led to little variation in results across metrics, with our selected prompt achieving more balanced results across metrics. For few-shot sampling, larger sets of 10 examples provided minor improvements on some metrics, but degradation in terms of compliance. Results with alternative zero-shot and few-shot prompting are provided in Appendix C.

### 4.3 Evaluation

Following Kim et al. (2021), we used three separate metrics to evaluate our models. SARI (Xu et al., 2016), the standard simplification metric as computed by the EASSE toolkit (Alva-Manchego et al., 2019),<sup>5</sup> which provides the proportion of n-grams that are added (*add* in what follows), kept (*keep*) or deleted (*del*). As a second metric, we used the normalised BertScore (hereafter, BESC) (Zhang et al., 2020b),<sup>6</sup> which has proved successful to measure meaning preservation. BLEU scores, computed with sacreBLEU (Post, 2018), are reported as well, as a precision measure of n-gram overlap between hypotheses and references.

We also included an additional metric for the task to indicate output compliance (CPL), i.e. the percentage of output sequences that include at least two separate sentences, to directly identify unsplit output. This metric was computed with a simple approximation, counting the number of end of sentence punctuation marks in the output sentence period, question and exclamation marks.

We compare our results with those obtained using the BiSECT and ABCD approaches of Kim

Model	Test	BESC	SARI	CPL
PRT-DE	DeSSE	86.81	61.02	71.03
PTG-DE		87.67	63.29	72.82
LoRA-DE		<b>88.78</b>	<b>69.47</b>	<b>94.64</b>
PRT-DE	BiSECT	73.42	27.03	63.10
PTG-DE		<b>73.77</b>	28.62	76.55
LoRA-DE		72.79	<b>30.85</b>	<b>93.45</b>
PRT-BI	DeSSE	81.95	37.42	6.55
PTG-BI		<b>82.00</b>	40.20	15.67
LoRA-BI		80.54	<b>48.80</b>	<b>71.43</b>
PRT-BI	BiSECT	37.00	1.55	8.79
PTG-BI		74.71	25.32	22.07
LoRA-BI		<b>76.12</b>	<b>38.93</b>	<b>81.55</b>

Table 2: Results for Prompt-tuning (PRT), P-Tuning (PTG) and LoRA adaptation with the Pythia-6.9B model trained on DeSSE (-DE) or BiSECT (-BI) data.

et al. (2021) and Gao et al. (2021), respectively, the state of the art for the selected corpora. For BiSECT, we used their publicly shared output directly, filtering the sentences that were excluded from the *Selected* test set described in Section 4.1, to maintain strict comparability with the other tested methods.<sup>7</sup> For ABCD, the output was computed with the publicly shared models and code,<sup>8</sup> on the *Selected* test subset, for a fair comparison.

## 5 SPRP with Large Language Models

### 5.1 Prompting Variants

Table 2 shows the results of soft-prompting and LoRA fine-tuning. Excepting two BertScore results, where P-Tuning was marginally better, LoRA adaptation with our hand-crafted simple prompt achieved the best results overall. The differences were relatively small between the three approaches in terms of SARI and BESC, except for PRT-BI on BiSECT with significant losses on all metrics.

The most notable result are the significantly larger differences in terms of compliance, where LoRA was markedly better in all cases and both soft-prompting methods failed significantly overall, drastically so when trained on BiSECT data. The poor performance of both methods on the more complex SPRP data is noteworthy, considering the gains achieved with soft-prompting in other downstream tasks. Considering these results, we selected

<sup>5</sup><https://github.com/feralvam/easse>

<sup>6</sup>[https://github.com/Tiiiger/bert\\_score](https://github.com/Tiiiger/bert_score)

<sup>7</sup><https://github.com/mounicam/BiSECT/tree/main/outputs/>

<sup>8</sup><https://github.com/serenayj/DeSSE>




Figure 1: Impact of parameter size on the DeSSE (left) and BiSECT (right) test sets, with Pythia model variants fine-tuned over data from the DeSSE (-DE) or BiSECT (-BI) training data.




Figure 2: Split and Rephrase results on the DeSSE (left) and BiSECT (right) test sets, for instruction-tuned Alpaca and Dolly model variants. Models were applied either in a zero-shot fashion (-ZS), or in a few-shot manner (-FS) with 5 examples sampled from the DeSSE (-DE) or BiSECT (-BI) datasets.

the more robust LoRA approach as our main adaptation method in the remaining experiments.

## 5.2 Fine-tuned LLM Variants

Figure 1 shows the results for the Pythia variants fine-tuned on DeSSE (-DE) or BiSECT (-BI).

In terms of parameter size, for in-domain scenarios, on all metrics there is a tendency towards increased accuracy as the number of parameters increases, although only marginally so and with both SARI and CPL actually incurring minor losses with the largest DeSSE-tuned 12B models. Absolute scores were lower with the BiSECT-tuned model on its own test set, which is not unexpected considering the complexity of the test data in this case. CPL and SARI for the BiSECT-tuned model were also more erratic, with a marked drop for the 1B model in particular, which could be due to the challenges posed by training on the complex BiSECT datasets for smaller models.

Domain shift results provided interesting insights regarding robustness of the variants. For the BiSECT-tuned model, switching from its in-

domain test set to the DeSSE test led to higher absolute scores for all metrics but CPL. This is likely due to the lesser complexity of the latter test, although the marked CPL drop could also indicate that training on the more complex BiSECT data resulted in limited portability for the critical splitting part of the task.

For DeSSE-tuned models, switching to the more complex BiSECT test did result in absolute drops on most metrics, as expected with a significant domain change. However, they achieved comparable results on CPL despite domains switch, and also achieved comparable results to the BiSECT-tuned model in terms of BESC, BLEU and SARI, while also largely outperformed it on CPL. Modelling SPRP on shorter sentences and limited domain variation, as featured in DeSSE, thus seem to be a better option for higher portability across domains.

## 5.3 Instruction-tuned Models

The zero-shot and in-context few-shot results with the Alpaca and Dolly models on the SPRP task are shown in Figure 2. The first notable result is the

Train	Size	BESC	BLEU	SARI	CPL
BI	1k	70.69	35.90	21.78	10.34
BI	2k	74.31	36.80	21.23	5.17
BI	3k	75.61	40.00	32.43	52.93
BI	4k	75.82	40.70	35.69	66.03
BI	6k	76.12	41.30	38.93	81.55
BI	14k	76.04	41.60	41.93	90.69
BI	30k	76.32	42.40	<b>45.06</b>	<b>95.17</b>
BI	62k	<b>76.74</b>	<b>42.60</b>	43.84	93.79
DE	1k	69.21	48.60	32.73	13.89
DE	2k	82.26	73.90	59.99	73.21
DE	3k	88.25	81.10	67.33	85.52
DE	4k	<b>88.80</b>	<b>81.60</b>	68.58	89.09
DE	6k	88.78	81.80	<b>70.39</b>	<b>94.64</b>

Table 3: Impact of training data size with the Pythia-6.9B model on BiSECT (BI) and DeSSE (DE)

overall higher scores obtained with Alpaca models, on all metrics, when compared to equivalent Dolly variants. The differences were not marginal as they amounted to average differences, between comparable models, of 13.06, 18.57 and 13.93 points in terms of BertScore, BLEU and SARI, respectively, on the DeSSE test set. On the BiSECT test, the average differences between comparable models amounted to 8.09, 7.23 and -0.97, with only the SARI metric leading to comparable results on this test set. The choice of a specific pre-trained instruction-tuned LLM for the SPRP task thus seems of particular importance, despite their relative similarity in terms of model architecture, training objectives and dataset variety.

Few-shot variants outperformed their zero-shot alternatives across the board, although it is worth noting that the Alpaca zero-shot model outperformed the Dolly few-shot model overall. As was the case for the fine-tuned Pythia variants, the results on the BiSECT test set were also significantly lower overall than on the DeSSE test, due in part to the higher complexity of this specific test set, which features significantly longer sentences on both the complex and split sides.

#### 5.4 Impact of Training Data Size

The results described in the previous sections for the best performing models were obtained with approximately 6,000 training data points. In this section, we aim to evaluate LLM accuracy on the SPRP task with either larger volumes of training

data, testing the ceiling of the approach, or smaller amounts to estimate minimal volumes of training data for this method. We used the fine-tuned *Pythia-6.9B* models for these experiments, per the previously discussed rationale.

The results are shown in Table 3. In both cases, scores increased across metrics as training size increased, although by smaller increments from the 3K mark upwards. Best results are split between the two largest training datasets on both BiSECT and DeSSE, indicative of a potential ceiling in terms of improvements from training data augmentation. It might also be the case that the observed plateauing resulted from a lack of variety in the additional training data, but verifying this hypothesis was beyond the scope of this work.

The two domains displayed noteworthy differences. Although using all the available training data resulted in better SARI and CPL scores on DeSSE, those obtained with 3K were closer to the best results than to those obtained with lower amounts. For BiSECT, the largest increase in terms of compliance was achieved at the 6K mark, with a stabilisation above 90% with training datasets of 14K or more samples. These results are in line with the previously discussed complexity of the BiSECT datasets, compared to DeSSE. It is worth noting that no model variant achieved compliance above 95.2, irrespective of training data size.

## 6 Comparative Evaluations

### 6.1 Metrics Results

From the results in the previous sections, we selected a subset of models for comparison with the ABCD and BiSECT approaches, namely: the best 6.9B Pythia models in terms of CPL and SARI (Section 5.4), which obtained only slightly lower results on all metrics than the 12B variant and features a comparable amount of parameters to the instruction-tuned LLMs; the Alpaca zero-shot model, representative of scenarios where no data are available for fine-tuning or in-learning; and the Alpaca few-shot model using either BiSECT and DeSSE examples. For the ABCD approach, we used both the bilinear and MLP variants.

We also included results from GPT-4 Turbo<sup>9</sup>, as one of the largest LLMs to date, although these results should be considered with caution: training

<sup>9</sup>Available at <https://help.openai.com/en/articles/8555510-gpt-4-turbo>. Version *gpt-4-1106-preview* February 2024. We used the same instruction as for our other zero-shot results.

Method	Test	BESC	BLEU	SARI	add	keep	del	CPL
pythia-6.9B-DE (6K)	DeSSE	<b>88.78</b>	<b>81.80</b>	<b>69.47</b>	<b>45.75</b>	<b>92.26</b>	<b>70.39</b>	94.64
alpaca-7B-ZS	DeSSE	75.46	64.60	56.44	23.14	89.35	56.82	83.93
alpaca-7B-FS-DE	DeSSE	84.55	74.90	56.08	27.76	88.92	51.56	68.65
ABCD-bilinear	DeSSE	66.27	57.20	52.54	17.63	81.04	58.96	<b>98.41</b>
ABCD-mlp	DeSSE	78.68	67.40	57.67	22.04	87.90	63.06	89.09
<i>GPT-4 Turbo</i>	DeSSE	<i>84.59</i>	<i>70.90</i>	<i>63.69</i>	<i>35.24</i>	<i>89.76</i>	<i>66.09</i>	<i>99.20</i>
pythia-6.9B-BI (6K)	BiSECT	76.12	41.30	38.93	12.92	63.59	40.27	81.55
pythia-6.9B-BI (30k)	BiSECT	<b>76.32</b>	<b>42.40</b>	<b>45.06</b>	<b>16.81</b>	<b>65.93</b>	<b>52.45</b>	95.17
alpaca-7B-ZS	BiSECT	69.32	36.20	29.82	5.08	60.39	23.99	88.62
alpaca-7B-FS-BI	BiSECT	73.17	36.80	35.82	8.57	60.46	38.43	65.86
BiSECT seq2seq	BiSECT	72.51	36.80	33.21	7.34	59.91	32.40	<b>99.48</b>
<i>GPT-4 Turbo</i>	BiSECT	<i>72.04</i>	<i>34.70</i>	<i>37.18</i>	<i>9.65</i>	<i>60.33</i>	<i>41.55</i>	<i>100.00</i>

Table 4: Comparative metrics results for the Split and Rephrase task on the DeSSE and BiSECT test sets.

data are unknown, with possible test data leakage, and reproducibility standards are not met, considering the proprietary nature of the model. The comparative results are shown in Table 4.

On both test sets, the fine-tuned Pythia models with 6.9B parameters outperformed all alternatives. On DeSSE, the best ABCD model (ABCD-mlp) trailed by 10.1, 14.4, 11.8 and 5.55 points in terms of BESC, BLEU, SARI and CPL, respectively; the ABCD-bilinear model obtained the best compliance score but the lowest results on the other metrics. On BiSECT, the BiSECT model was outperformed by 3.81, 5.6 and 11.85 points on BESC, BLEU and SARI, respectively; it was however the most compliant model. The LLM variants can thus provide more accurate results across the board, but can still lag behind alternatives in terms of enforcing split SPRP hypotheses.

It is worth noting that instruction-tuned models prompted with only 5 in-context examples obtained competitive results across the board, although with significantly lower compliance scores overall. The zero-shot results were also competitive against the ABCD and BiSECT approaches, although they obtained the lowest scores overall. The fine-tuned Pythia models were thus shown to be a better alternative when training datasets of sufficient size are available, but the instruction-based models could provide a reasonable alternative.

Finally, although GPT-4 Turbo achieved the highest compliance, it was outperformed by our best LoRA variants on all other metrics. After manual examination of sampled output, the main reason seems to be a stronger tendency to rephrase

and extend the original content. Different prompts might mitigate this issue, an option we left aside to maintain a fair comparison between models.

## 6.2 Manual Ranking Results

Metrics such as BLEU have important limitations for the SPRP task, displaying low correlation with human judgements in some experiments (Zhang et al., 2020a). Although the results of the previous section show large improvements across multiple metrics, indicative of marked improvements achieved by LLM-based SPRP, we conducted an additional manual comparative evaluation. For this specific evaluation, human evaluators were presented a source sentence and two SPRP hypotheses, which they had to rank in a 3-way evaluation, with either one of the hypotheses considered better, or both considered equally correct or incorrect.

For this task, we adapted the Appraise environment (Federmann, 2018), typically used for machine translation evaluation campaigns, and asked 11 evaluators to perform two sub-tasks, each involving 100 samples randomly selected: compare hypotheses from fine-tuned Pythia-6.9B models with those of ABCD-mlp and BiSECT models, on the DeSSE and BiSECT test sets, respectively. Evaluators were asked to compare the hypotheses according to SPRP criteria of compliance, sensuality, grammatically, missed facts and omitted facts.<sup>10</sup> The results are shown in Table 5.

Compared to ABCD, LLM-based hypotheses were largely preferred by human evaluators, at 70.27%, with 28.55% considered equal. On the

<sup>10</sup>See Appendix E for details on the experimental protocol.

S	P > S	P = S	P < S	$\alpha$
ABCD	70.27%	28.55%	1.18%	0.52
BiSECT	43.09%	38.45%	18.45%	0.49

Table 5: 3-way ranking results comparing hypotheses from Pythia-6.9B-DE (P) and alternative SPRP methods (S) ABCD or BiSECT, on the DeSSE and BiSECT test sets, respectively.  $\alpha$  indicates Krippendorf’s inter-annotator agreement.

more complex BiSECT test data, which mainly featured long sentences in the pharmaceutical and medical domain, the results were more balanced, with 38.45% of hypotheses considered equally good or bad. However, LLM hypotheses were still markedly preferred for the remaining hypotheses, at 43.09% against 18.45% of hypotheses selected from the BiSECT sequence-to-sequence model. In both evaluations, inter-annotator agreement was moderate overall, at 0.52 and 0.49, slightly lower on the more complex BiSECT data.

The results obtained in this comparative human evaluation are thus in line with those obtained on automated metrics overall, with marked improvements achieved via LLM-based SPRP.

## 7 Qualitative Evaluation

To gain further insights on the SPRP capabilities of LLMs, we performed an additional human evaluation, this time centred on assessing the validity of individual SPRP hypotheses on a set of predefined criteria. For this qualitative evaluation, we sampled 100 source sentences from Wiki-BM (Zhang et al., 2020a), a manually crafted benchmark created from Wikipedia data, and generated SPRP hypotheses with the Pythia-6.9B model fine-tuned on DeSSE data. The evaluation involved 8 human evaluators, who had to answer 6 questions related to the quality of an SPRP hypothesis considering a source sentence, addressing sencicality, grammaticality, deletion or introduction of new facts, correct splits and insufficient splits.<sup>11</sup>

Overall, the results, shown in Table 6, were largely positive. Most output was judged sensical, grammatical, with almost no introduction of new facts, and correct splits in over 91% of the cases. All were judged with almost perfect or strong inter-annotator agreement, with sencicality and correct

<sup>11</sup>The questions were adapted from Zhang et al. (2020a), with minor modifications. See Appendix F for details on the experimental protocol.

Question	Answer	$\alpha$
Sensical? [0-2]	86.12% / 1.84	0.803
Grammatical? [0-2]	92.62% / 1.91	0.907
No missed facts? [Y/N]	84.13%	0.597
No new facts? [Y/N]	99.50%	0.996
Correct split? [Y/N]	91.13%	0.797
Enough split? [Y/N]	83.25%	0.548

Table 6: Qualitative evaluation results with Pythia-6.9B on Wiki-BM ( $\alpha$  = Krippendorf’s alpha inter-rater agreement). Answers for the first two questions are the percentage of 2 and the average score on a [0-2] scale. For the other questions, we report the percentage of Yes.

splitting slightly below at around 80% agreement.

Although still featuring scores above 80%, the remaining aspects led to comparatively lower marks, with missed facts and not enough splits in 15.87% and 16.75% of the cases, respectively, and inter-rater agreement between 0.55 and 0.60. Missed facts in particular were somewhat more difficult to judge, according to post-mortem feedback, as some rephrasing cases involved implicit rather than explicit logical connection between sentences, which could be rated more or less strictly. Similarly, the need for additional splits can be viewed as more subjective than other categories, as evaluators could view some grammatically correct additional splits as cases of over-splitting.<sup>12</sup>

It is worth noting that these results were obtained with a model fine-tuned on data from a different domain than Wiki-BM, namely DeSSE, showing that knowledge extracted from one dataset could be effectively transferred to other domains.

## 8 Conclusions

In this work, we evaluated the performance of large language models on the Split and Rephrase task, improving over the state of the art by significant margins on all metrics over SPRP datasets of varying complexity. The results of a manual 3-way ranking evaluation further confirmed human preference for LLM-based hypotheses, and a qualitative evaluation on Wiki-BM data further demonstrated the quality of the generated SPRP output.

Among LLM variants, models fine-tuned with LoRA achieved the best results overall. Zero-shot and few-shot instruction-tuned variants achieved

<sup>12</sup>Examples of LLM output from this evaluation are shown in Appendix B. Metrics results on the complete Wiki-BM benchmark are indicated in Appendix D.



reasonable performance, modulo a significant drop in splitting compliance for the latter, and might thus constitute a reasonable alternative in languages where training corpora are lacking for the task.

Increasing fine-tuning data size resulted in only gradual improvements, plateauing on most metrics with subsets of the available data. Increased parameters led to only minor improvements, with models around 7B achieving the more balanced results. Despite significant improvements, none of the examined variants achieved full task compliance, and core aspects of SPRP remain a challenge.

## 9 Limitations

The first limitation of an LLM-based approach is the need for pretrained LLMs for the specific languages in which the SPRP task needs to be performed. Note however that even English-centric LLMs can exhibit multilingual capability (Armengol-Estapé et al., 2022), thus mitigating the issue. Additionally, multilingual models such as BLOOM (Scao et al., 2022) can provide support for the task; we provide preliminary results in this sense in Appendix G. A more in-depth evaluation of LLM performance on the SPRP task, for languages with limited coverage, would be required to further assess the limitations of the approach.

A second limitation concerns the dependency on training data to fine-tune pretrained LLMs. We established that around 3K pairs of complex and split sentences might be necessary for the task, with high quality results obtained using between 4K and 6K training pairs, depending on the domain. Although this type of training data can be collected from either parallel multilingual datasets, as in BiSECT, or specific types of essays, as in DeSSE, this still requires dedicated work for new languages. Machine translating existing SPRP corpora with high quality models could provide a basis for fine-tuning in other languages, although a precise evaluation of this approach would be needed (see Appendix G for preliminary results along these lines). We also showed that instruction-tuned pretrained models could provide quality results using only 5 in-context learning pairs, which could be rapidly crafted manually.

## 10 Ethical Considerations

Our approach involves the use of large language models, whose computational performance is typically higher when deployed in more powerful en-

vironments with GPUs, for either training or inference. Under such usage, electric consumption and associated carbon footprint are likely to increase and users of our method under these conditions should be aware of this type of impact. Our approach mitigates these costs via the use of pretrained language models of relatively limited size, and an efficient tuning method such as LoRA (Hu et al., 2022) which minimises training time and computational resources.

## Acknowledgements

We wish to thank the anonymous ARR reviewers for their helpful comments, and the participants in our manual evaluations for their time and contribution. This work was partially supported by the Department of Economic Development and Competitiveness of the Basque Government (Sprig Group) through funding for the IRAZ project (ZL-2024/00570).

## References

- Roei Aharoni and Yoav Goldberg. 2018. [Split and rephrase: Better evaluation and stronger baselines](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 719–724, Melbourne, Australia. Association for Computational Linguistics.
- AlMotasem Bellah AlAjilouni and Jinlong Li. 2023. Knowledge transfer to solve split and rephrase. In *2023 International Conference on Information Technology (ICIT)*, pages 680–685. IEEE.
- Fernando Alva-Manchego, Louis Martin, Carolina Scarton, and Lucia Specia. 2019. [EASSE: Easier automatic sentence simplification evaluation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 49–54, Hong Kong, China. Association for Computational Linguistics.
- Jordi Armengol-Estapé, Ona de Gibert Bonet, and Maite Melero. 2022. [On the multilingual capabilities of very large-scale English language models](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3056–3068, Marseille, France. European Language Resources Association.
- Paulo Berlanga Neto and Evandro Eduardo Seron Ruiz. 2021. [Split-and-rephrase in a cross-lingual manner: A complete pipeline](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 155–164, Held Online. INCOMA Ltd.

- Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. 2023. [Pythia: A suite for analyzing large language models across training and scaling](#).
- Jan A. Botha, Manaal Faruqui, John Alex, Jason Baldridge, and Dipanjan Das. 2018. [Learning to split and rephrase from Wikipedia edit history](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 732–737, Brussels, Belgium. 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, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- Bidyut Das, Mukta Majumder, and Santanu Phadikar. 2018. A novel system for generating simple sentences from complex and compound sentences. *International Journal of Modern Education and Computer Science*, 12(1):57.
- 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.
- Christian Federmann. 2018. Appraise evaluation framework for machine translation. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 86–88.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. The Pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Yanjun Gao, Ting-Hao Huang, and Rebecca J. Passonneau. 2021. [ABCD: A graph framework to convert complex sentences to a covering set of simple sentences](#). 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 3919–3931, Online. Association for Computational Linguistics.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640, Berlin, Germany. Association for Computational Linguistics.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Joongwon Kim, Mounica Maddela, Reno Kriz, Wei Xu, and Chris Callison-Burch. 2021. [BiSECT: Learning to split and rephrase sentences with bitexts](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6193–6209, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2023. Gpt understands, too. *AI Open*.
- Eric Malmi, Sebastian Krause, Sascha Rothe, Daniil Mirylenka, and Aliaksei Severyn. 2019. [Encode, tag, realize: High-precision text editing](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5054–5065, Hong Kong, China. Association for Computational Linguistics.
- Shashi Narayan and Claire Gardent. 2014. [Hybrid simplification using deep semantics and machine translation](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 435–445, Baltimore, Maryland. Association for Computational Linguistics.
- Shashi Narayan, Claire Gardent, Shay B. Cohen, and Anastasia Shimorina. 2017. [Split and rephrase](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 606–616, Copenhagen, Denmark. Association for Computational Linguistics.
- Christina Niklaus, Matthias Cetto, André Freitas, and Siegfried Handschuh. 2019. [Transforming complex sentences into a semantic hierarchy](#). In *Proceedings of the 57th Annual Meeting of the Association for*

- Computational Linguistics*, pages 3415–3427, Florence, Italy. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
- Matthew Shardlow. 2014. A survey of automated text simplification. *International Journal of Advanced Computer Science and Applications*, 4(1):58–70.
- Felix Stahlberg and Shankar Kumar. 2020. [Seq2Edits: Sequence transduction using span-level edit operations](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5147–5159, Online. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. [Pointer networks](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. [Optimizing statistical machine translation for text simplification](#). *Transactions of the Association for Computational Linguistics*, 4:401–415.
- Li Zhang, Huaiyu Zhu, Siddhartha Brahma, and Yunyao Li. 2020a. [Small but mighty: New benchmarks for split and rephrase](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1198–1205, Online. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Zhemin Zhu, Delphine Bernhard, and Iryna Gurevych. 2010. [A monolingual tree-based translation model for sentence simplification](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 1353–1361, Beijing, China. Coling 2010 Organizing Committee.

Original Text	Split & Rephrased Text
Most people travel to Europe or the Islands and get a taste of the culture in those places.	Most people travel to Europe or the Islands . They get a taste of the culture in those places.
People would've been killed , places would have been bombed and there would have been so much going on all at once.	People would've been killed . Places would have been bombed . There would have been so much going on all at once.
War is injuring innocent people which is not necessary.	War is injuring innocent people . It is not necessary.
As education continues to reach the poorer areas of our world i truly believe that one day we will all live in TRUE harmony.	Education continues to reach the poorer areas of our world . I truly believe that one day we will all live in TRUE harmony.
This land is still the home to many civilians who rely on the resources in their area to live.	This land is still the home to many civilians . They rely on the resources in their area to live.

Table 7: SPRP examples on the DeSSE dataset

## A LLM Training and Inference

Fine-tuning was performed with LoRA (Hu et al., 2022) using the following default parameters:  $r=8$ ,  $\alpha=16$ ,  $\text{dropout}=0.05$ ; targeted modules were: Query, Key and Value. We also used default generation parameters, namely:  $\text{temperature}=0.1$ ;  $\text{top}_p=0.75$ ;  $\text{top}_k=40$ ;  $\text{num\_beams}=4$ ;  $\text{max\_new\_tokens}=128$ .

Each training sample was provided in the following format, where  $\alpha$  is a complex sentence and  $\beta$  its corresponding split sentences, sampled from the DeSSE and BiSECT training sets:

INSTRUCTION: Split the sentence below into separate sentences.

INPUT:  $\alpha$

RESPONSE:  $\beta$

The following text was prepended to each sample prompt: *"Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request."*

Each training process was performed on 1 A40 GPU with 48GB of RAM. As reference times, fine-tuning took 06:07:59 hours for the Pythia-6.9B model on the BiSECT 6K training dataset and 48 hours over the larger 64K dataset. Fine-tuning the Pythia 12B model took 06:07:59 hours and the smaller 410M model took 00:37:51 hours, on the 6K BiSECT dataset.

For soft-prompting methods, we used the following default settings:  $\text{num\_virtual\_tokens}=20$ ,  $\text{prompt\_tuning\_init}=\text{"Split the sentence below into separate sentences."}$  for Prompt-tuning and  $\text{num\_virtual\_tokens}=20$ ,  $\text{encoder\_hidden\_size}=128$  for P-tuning.

## B Split and Rephrase Examples

In Tables 7, 8 and 9, we provide illustrative examples of split and rephrased output obtained with the fine-tuned Pythia-6.9B model on Dese, BiSECT and Wiki-BM source sentences, respectively. We selected examples that are representative of the type of alterations to the source performed by the selected model, which was used for the manual evaluations, including errors in splitting and/or rephrasing.

Sequences from the source text that have been correctly deleted or modified are shown in purple. Sequences that have been correctly added to the output are shown in green. Sequences that have been incorrectly altered in the original sentence are shown in red.

Original Text	Split & Rephrased Text
Dialysis does not alter serum doxycycline half-life and thus would not be of benefit in treating cases of overdose.	Dialysis does not change the serum half-life of doxycycline. Thus, it would not be of benefit in the treatment of cases of overdose.
In the first study, type A resistant patients (A-non responders) were randomised to receive placebo or 10, 000 U of NeuroBloc and in the second, type A toxin responsive patients (A-responders) were randomised to receive placebo, 5000 U or 10, 000 U of toxin.	In the first study, type A non-responders were randomized to receive placebo or 10, 000 U of NeuroBloc. In the second study, type A responders were randomized to receive placebo, 5000 U or 10, 000 U of toxin.
Most patients treated with Zevalin experience a severely reduced number of blood cells, which needs some time to normalise again after treatment.	Most patients treated with Zevalin experience a severe reduction in the number of blood cells. It takes some time to return to normal after treatment.
Lisonorm tablet contains 5mg amlodipine / 10mg lisinopril and is an anti-hypertensive indicated in patients with blood pressure adequately controlled with lisinopril and amlodipine given concurrently at the same dose level.	Lisonorm tablet contains 5 mg of amlodipine and 10 mg of lisinopril. It is indicated for the treatment of hypertension in patients whose blood pressure is adequately controlled with lisinopril and amlodipine.

Table 8: SPRP examples on the BiSECT dataset

Original Text	Split & Rephrased Text
She walked out on their wedding, and she currently is single and holds intense anger towards Richard for ruining their wedding, and also became furious to learn she is pregnant, even though Richard told her he got a vasectomy.	She walked out on their wedding. She currently is single. She holds intense anger towards Richard for ruining their wedding. She became furious to learn she is pregnant. Richard told her he got a vasectomy.
Southfort is an unincorporated community in DeSoto County, Florida, United States, located approximately south of Fort Ogden on U.S. Route 17.	Southfort is an unincorporated community in DeSoto County, Florida. It is located approximately south of Fort Ogden on U.S. Route 17.
Parke played for powerhouse youth club FC Delco, and played four years of college soccer at Drexel University, before being drafted by the MetroStars with 60th and last overall pick of the 2004 MLS SuperDraft.	Parke played for powerhouse youth club FC Delco. He played four years of college soccer at Drexel University. He was drafted by the MetroStars with 60th and last overall pick of the 2004 MLS SuperDraft.
Auguste Le Duc started on a bridge between the two islands, which was quickly swept away by the forces of nature.	Auguste Le Duc started on a bridge between the two islands. The bridge was quickly swept away by the forces of nature.

Table 9: SPRP examples on the Wiki-BM dataset

Test	Method	BESC	BLEU	SARI	add	keep	del	CPL
DeSSE	5-shot-DE	92.51	74.90	56.08	27.76	<b>88.92</b>	51.56	68.65
	10-shot-DE	<b>92.75</b>	<b>75.10</b>	<b>58.07</b>	<b>31.49</b>	88.46	<b>54.26</b>	<b>71.83</b>
	5-shot-BI	88.04	51.00	42.53	11.51	73.53	42.55	31.94
	10-shot-BI	87.40	47.90	40.09	8.55	71.44	40.28	21.63
BiSECT	5-shot-BI	86.99	36.80	35.82	8.57	<b>60.46</b>	38.43	65.86
	10-shot-BI	86.60	35.70	<b>38.06</b>	<b>10.08</b>	60.03	<b>44.09</b>	64.31
	5-shot-DE	87.63	<b>38.50</b>	28.76	5.46	59.92	20.89	66.38
	10-shot-DE	<b>87.73</b>	38.00	29.32	5.74	59.62	22.60	<b>68.97</b>

Table 10: Extended few-shot results with the Alpaca-7B model using examples from DeSSE (DE) and BiSECT (BI).

## C Extended Prompting Variants

**Few-shot prompts.** As noted in Section 4.2, we explored the use of larger few-shot sampling on the instruction-tuned models, using 10 in-context learning examples instead of the 5 selected for our core experiments. Results of these additional experiments are presented in Table 10 for the best-performing instruction-tuned Alpaca 7B model.

Prompting with 10 examples led to relatively minor gains in SARI in most cases, with a marked tendency to delete more content in the rephrased version. BertScore and BLEU were virtually similar. Extending the in-context samples also resulted in CPL losses with domain shift, up to 10 points with 10-shot-BI on the DeSSE test set, notably, although the reverse was observed with 10-shot-DE on the BiSECT test, with a 2.59 points gain. A minor CPL loss was also observed with 10-shot-BI over BiSECT.

Domain shift losses were not totally surprising considering the in-context conditioning on samples from the very different DeSSE and BiSECT domains. Considering these results, in particular CPL losses, we selected 5-shot as our primary method for few-shot in-context learning.

**Zero-shot prompts.** The zero-shot prompt selected for our experiments was chosen after preliminary testing of different manually engineered variants. No clear better prompting candidate emerged from this initial phase, although it is always possible that an alternative prompt could lead to better result with a given model. To provide further information on zero-shot prompting variants, in Table 11 we report additional results with the following prompts, against GPT-4 Turbo (*gpt-4-1106-preview*, accessed February 2024):

- Prompt O: "Split the sentence below into separate sentences."
- Prompt A: "Split the following sentence into multiple shorter sentences, while preserving the original meaning as closely as possible."
- Prompt B: "Rephrase the given sentence to make it more concise and direct, without changing the core meaning."
- Prompt C: "Split the following sentence into multiple concise sentences, separating each distinct idea."

Prompt O is our original prompt, used throughout our experiments, and prompts A-C are alternative manually engineered prompts for the task.<sup>13</sup> Leaving aside prompt B, which resulted in almost no compliance and severe degradation overall, the other two prompt variants achieved relatively similar results compared to Prompt O. Overall, the latter still achieved better results on all metrics on DeSSE, and all but SARI on BiSECT. All prompt variants were outperformed by the best fine-tuned 6.9B variants, except on compliance.

<sup>13</sup>We thank the anonymous ARR reviewer who suggested these specific prompt variants as additional test cases.

Method	Test	BESC	BLEU	SARI	add	keep	del	CPL
GPT-4 Prompt O	DeSSE	<b>84.59</b>	<b>70.90</b>	<b>63.69</b>	<b>35.24</b>	<b>89.76</b>	<b>66.09</b>	99.20
GPT-4 Prompt A	DeSSE	79.92	59.40	56.52	25.74	82.94	60.86	99.80
GPT-4 Prompt B	DeSSE	61.87	15.40	26.27	2.58	37.18	39.05	1.38
GPT-4 Prompt C	DeSSE	79.27	58.50	55.45	24.24	81.89	60.24	<b>100.00</b>
<i>pythia-6.9B-DE (6K)</i>	DeSSE	88.78	<i>81.80</i>	<i>69.47</i>	<i>45.75</i>	<i>92.26</i>	<i>70.39</i>	<i>94.64</i>
GPT-4 Prompt O	BiSECT	<b>72.04</b>	<b>34.70</b>	37.18	9.65	60.33	41.55	<b>100.00</b>
GPT-4 Prompt A	BiSECT	67.78	29.30	42.54	10.36	59.92	<b>57.33</b>	<b>100.00</b>
GPT-4 Prompt B	BiSECT	65.37	18.40	<b>43.38</b>	9.04	49.90	71.20	3.96
GPT-4 Prompt C	BiSECT	68.14	30.60	42.59	<b>10.46</b>	<b>60.61</b>	56.69	<b>100.00</b>
<i>pythia-6.9B-BI (30K)</i>	BiSECT	76.32	<i>42.40</i>	<i>45.06</i>	<i>16.81</i>	<i>65.93</i>	<i>52.45</i>	<i>95.17</i>

Table 11: Zero-shot prompt variant results on DeSSE and BiSECT with GPT4. Pythia-6.9B results are provided as a reference. Best results amongst GPT-4 variants are shown in bold.

Method	BLEU	SARI	add	keep	del	BESC	CPL
pythia-6.9B-DE	72.90	65.10	36.14	91.15	68.10	85.50	98.19
<i>T<sub>MWS</sub></i> (AlAjlouni and Li, 2023)	59.50	60.00	-	-	-	-	-
<i>GPT-4 Turbo</i>	74.40	68.63	39.96	91.90	74.03	86.23	99.70

Table 12: Metrics results on the Wiki-BM benchmark test set. Comparison between Pythia-6.9B fine-tuned on DeSSE data, results reported in (AlAjlouni and Li, 2023), and GPT4 Turbo (accessed on February 2024).

## D Evaluation Metrics on Wiki-BM

In section 7, we used portions of the Wiki-BM corpus (Zhang et al., 2020a) for our qualitative evaluation of LLM-based SPRP output. In Table 12, we provide the results in terms of metrics for the complete Wiki-BM test set, along with results from AlAjlouni and Li (2023), as representative of an alternative method based on fine-tuning an encoder-decoder pretrained T5 model, and GPT4 Turbo, as one of the largest LLMs to date. The caveats mentioned in Section 6 regarding the latter apply to these results as well.

The results for our LLM-based approach are in line with those obtained on the DeSSE test set across metrics: slightly lower on BESC, BLEU, and SARI, but higher in terms of compliance (CPL). The Pythia-6.9B model fine-tuned with LoRA notably outperformed the reported scores in AlAjlouni and Li (2023) by 13.4 and 5.1 points on the BLEU and SARI metrics, respectively. GPT-4 Turbo achieved the best results on all metrics, but by relatively small margins overall. These results confirm the viability of relatively smaller LLMs, tuned with LoRA on small amounts of data, for the SPRP task.

## E 3-Way Ranking Evaluation Protocol

The 3-way ranking manual evaluation involved 11 human evaluators, all researchers with at least upper-intermediate knowledge of English, who were not paid for the task. The evaluation was split into two tasks: for one task, the evaluators were asked to rank the output of the fine-tuned Pythia-6.9B model and the output of the ABCD model; for the other task, they had to rank the output of the fine-tuned Pythia-6.9B model and the output of the BiSECT model. Each of the two tasks involved 100 samples. Evaluators were also asked to ignore minor capitalisation and tokenisation issues, which were typically present in the output of the ABCD and BiSECT models, as detected in preliminary data inspection.

We adapted the Appraise framework (Federmann, 2018) for the task, mainly by simplifying the presentation of the options, and the following instructions were provided to complete the task:

- The task is named Split and Rephrase (SPRP) and consists in splitting a single complex sentence into a sequence of two or more grammatical sentences and rephrasing them as needed to preserve the

It shows how innocent these people are and that war should never be a solution.  
— Source

**Better** It shows how innocent these people are. War should never be a solution.  
— Translation A

**Equal**

**Better** It shows how innocent these people are. innocent people war should never be a solution.  
— Translation B

**Flag Error**

Figure 3: Appraise interface adapted for the 3-way ranking SPRP evaluation

original meaning without omitting or adding information.

- You will be shown 100 examples for each task, each example as shown in Figure 3, where:
  - The sentence shown at the top (source) is the original single sentence.
  - The sentences shown as translation A and translation B are SPRP hypothesis which you need to compare.
  - After reading the source sentence and each SPRP hypothesis, you should perform one of the following actions:
    - \* Click on the *Better* button next to one of the hypothesis if it is better than the other hypothesis according to the criteria below;
    - \* Click on the *Equal* button if the two hypothesis are both equally correct or both equally incorrect, according to the criteria indicated below.
- An SPRP hypothesis can be considered correct if:
  - it contains at least two separate sentences (*compliance*),
  - it contains all the information in the source sentence (*no missing facts*),
  - it does not contain information that was not present in the source sentence (*no new facts*),
  - all the sentences in the hypothesis are grammatical (*grammaticality*),
  - all the sentences in the hypothesis make sense both in isolation and in sequence (*sensicality*).
- Ranking criteria:
  - If both hypotheses meet all criteria, select *Equal*.
  - If neither hypothesis meets any of the criteria, select *Equal*.
  - If one hypothesis meets all criteria and the other does not, select the former as *Better*.
  - If one of the hypotheses is markedly better than the other, in the sense that as it meets almost all criteria whereas the other fails to meet most criteria, and only contains minor errors that do not impede comprehension of the original information, you may select the former as *Better*. In case of doubt, select *Equal*.
  - When validating an hypothesis, please ignore minor capitalisation errors (e.g., at the beginning of a sentence) and minor tokenisation errors (e.g., punctuation preceded by a space).<sup>14</sup>

System hypotheses are presented in random order by default in the Appraise environment.



Guy Mathez is a Swiss football manager from Canton of Jura.
Guy Mathez is a Swiss football manager. He is from Canton of Jura.

1. Is it sensical?

0  1  2

2. Is it grammatical?

0  1  2

3. Does it miss any existing facts?

Yes  No

4. Does it introduce new facts?

Yes  No

5. Does it have splits in the wrong place or no split at all?

Yes  No

6. Should some of its sentences be further split?

Yes  No

Figure 4: Web interface for the qualitative evaluation

## F Qualitative Evaluation Protocol

The evaluation involved 8 human evaluators, all researchers with at least upper-intermediate knowledge of English, who were not paid for the task. Each evaluator had to answer 6 questions related to the quality of an SPRP hypothesis considering a source sentence. We randomly sampled 100 sentences from the Wiki-BM corpus (Zhang et al., 2020a) and generated SPRP hypotheses with the Pythia-6.9B model fine-tuned on DeSSE data.

The questions were taken from Zhang et al. (2020a) for the most part, with minor modifications: a simplified 0-2 Likert scale instead of their 0-5 scale for the first two questions, as the latter appeared more difficult to use in practice, and a reformulation of question 5 to explicitly include cases where the output was not split.

The samples and associated questions were provided within an ad-hoc Web-based environment developed for the task. The following instructions were provided to the evaluators:

- Upon entering the Web-based evaluation environment, you will be shown a page as in Figure 4.
- The topmost sentence is the original source sentence. The sentence below is a Split & Rephrase (SPRP) hypothesis, i.e., a rewrite of the source sentence into a sequence of separate sentences that preserve the meaning of the source sentence.
- For each pair of source sentence and SPRP hypothesis, you will need to answer all 6 questions in the page. The guidelines for each question are indicated below:

### – Is it sensical? (scale of 0-2)

<sup>14</sup>This advice was meant to avoid discarding the ABCD or BiSECT hypotheses on these non-SPRP criteria, as they were shown to contain this type of errors in their respective output, upon preliminary data inspection.

- \* Indicate:
    - 0 if the SPRP hypothesis contains major sensicality errors (the text is barely comprehensible);
    - 1 if it contains minor sensicality errors (the text is comprehensible but has lost part of the original meaning, e.g., lacking logical connection between sentences);
    - 2 if there are no sensicality errors.
  - \* Examples:
    - Source: *"It was a great evening, until John arrived."*
    - Example of 0: *"It was an evening. The great John arrived."*
    - Example of 1: *"It was a great evening. John arrived."*
    - Example of 2: *"It was a great evening. Until John arrived."*
- Is it grammatical? (scale of 0-2)**
- \* Indicate:
    - 0 if the SPRP hypothesis contains major grammar errors, which impede text comprehension;
    - 1 if it contains minor grammatical errors, which barely impede text comprehension;
    - 2 if there are no grammar errors.
  - \* Examples:
    - Source: *"There was a treasure, which was buried on the island."*
    - Example of 0: *"There was a treasure which. He buried on the island."*
    - Example of 1: *"There was a treasure. It buried on the island."*
    - Example of 2: *"There was a treasure. It was buried on the island. "*
- Does it miss any existing fact? (yes/no)**
- \* Indicate:
    - *Yes* if the SPRP hypothesis omits information from the source sentence;
    - *No* if the SPRP hypothesis does not omit information from the source sentence.
  - \* Examples:
    - Source: *"There was a treasure on a Greek island, which was never discovered."*
    - Example of Yes: *"There was a treasure. It was never discovered."*
    - Example of No: *"There was a treasure on a Greek island. It was never discovered."*
- Does it introduce new facts? (yes/no)**
- \* Indicate:
    - *Yes* if the SPRP hypothesis includes information that was not present in the source sentence;
    - *No* if the SPRP hypothesis does not include information that was not present in the source sentence.
  - \* Examples:
    - Source: *"There was a treasure, which was never discovered."*
    - Example of Yes: *"There was a treasure on a British island. It was never discovered."*
    - Example of No: *"There was a treasure. It was never discovered."*
- Does it have splits in the wrong place or no split at all? (yes/no)**
- \* Indicate:
    - *Yes* if the SPRP hypothesis includes sentence splits in the wrong places or consists of only one sentence;
    - *No* if the SPRP hypothesis does not include sentence splits in the wrong places and contains at least two separate sentences;
  - \* Examples:
    - Source: *"There was a treasure on a Greek island, which was never discovered and remained there forever."*

Examples of Yes:

*"There was a treasure. On a Greek island it was never discovered. It remained there forever."*

*"There was a treasure on a Greek island, which was never discovered and remained there forever."*

Examples of No:

*"There was a treasure on a Greek island. It was never discovered. It remained there forever."*

*"There was a treasure on a Greek island. It was never discovered and remained there forever."*

**– Should some of its sentences be further split? (yes/no)**

\* Indicate:

- Yes if the SPRP hypothesis could be further split and rephrased into more sentences.
- No if the SPRP hypothesis could not be further split and rephrased into more sentences.

\* Examples:

Source: *"There was a treasure on a Greek island, which was never discovered and remained there forever."*

Examples of Yes:

*"There was a treasure on a Greek island. It was never discovered and remained there forever."*

*"There was a treasure on a Greek island, which was never discovered. It remained there forever."*

Example of No: *"There was a treasure on a Greek island. It was never discovered. It remained there forever."*

- NOTE: Some of the source sentences contain errors, of the types listed below. These cases should be evaluated as if the source sentence was a correct complex sentence:<sup>15</sup>
  - The source sentence actually contains two sentences, with missing space around the sentence separator, e.g., "xxx.They xxx".
  - There is a missing comma, e.g. "They are representatives from Tucson, Arizona who xxx".
  - There is a missing preposition, e.g. "The plan seemed work".

## G Initial Multilingual Results

Our results in this work were based on English, for which different LLM variants and datasets were available. As a first approximation in other languages, in this section we provide results for German, French and Spanish, using a multilingual BLOOM 7.1B model<sup>16</sup> (Scao et al., 2022) to ensure better coverage than with the English-centric Pythia model. We also include English results with the BLOOM model, for comparison purposes.

Since the DeSSE datasets were only available in English, we used an NLLB MT model<sup>17</sup> (Costa-jussà et al., 2022) to translate the corpora into each language. Although this approximation can introduce translation errors in the process, it was the only available alternative for the DeSSE scenario.

The BiSECT corpora were available for all three of the aforementioned languages, but suffered from key issues: to the exception of the French training corpus, all datasets were missing diacritics, which are critical in these languages. To partially cope with this issue, for French we created an additional test set by randomly extracting 2,000 pairs from the original training dataset; for the other languages, we report results using the faulty datasets, as there was no alternative in these cases. Train and dev size were the

<sup>15</sup>These cases had eluded our initial manual review of the data sampled from Wiki-BM. There were 10 such cases in total in the evaluation, and only 2 for the most problematic first case.

<sup>16</sup><https://huggingface.co/bigscience/bloom-7b1>

<sup>17</sup><https://huggingface.co/facebook/nllb-200-3.3B>

Language	Train	Test	BESC	BLEU	SARI	add	keep	del	CPL
English	DeSSE	DeSSE	89.05	82.00	69.56	45.51	92.42	70.74	94.05
English	BiSECT	BiSECT	76.18	41.50	42.10	14.90	64.45	46.95	81.72
French	DeSSE	DeSSE	89.30	65.00	53.52	24.79	81.18	54.59	83.93
French	BiSECT	BiSECT*	80.61	27.20	38.45	3.88	56.55	54.92	80.21
French	BiSECT	BiSECT	80.60	32.60	34.89	8.28	55.50	40.87	72.55
Spanish	DeSSE	DeSSE	90.67	67.10	54.30	23.95	83.28	55.67	87.70
Spanish	BiSECT*	BiSECT*	83.02	39.30	30.00	5.46	60.16	24.38	57.03
German	DeSSE	DeSSE	87.62	54.80	44.75	16.47	72.85	44.93	72.22
German	BiSECT*	BiSECT*	78.19	26.80	26.02	1.67	47.92	28.47	63.13

Table 13: Results in English, German, French and Spanish with the Bloom-7B1 model. Train and test datasets with missing diacritics are indicated with an asterisk. BertScore results are only normalised for English.

same as for English, whereas the test sets contained 1036, 2018 and 735 pairs for French, Spanish and German, respectively.

The results for these experiments are summarised in Table 13. The English results obtained with the BLOOM model were similar to those obtained with the Pythia model (see Table 4), indicating that the overall SPRP ability does not seem impacted by variance between fine-tuned models of similar overall architecture and parameter size.

On the machine-translated DeSSE datasets, the results were high overall across metrics for Spanish, French, German, in descending accuracy order, although below those obtained in English, presumably as a result of the comparatively smaller volumes of LLM training data in languages other than English. On BiSECT, of note are the results obtained with models trained and tested on datasets that lacked diacritics. These models still generated output that included correct word forms with diacritics, the underlying model having been trained on correct data. The scores obtained against test sets without diacritics were thus negatively affected, and all models but the French one were trained on faulty data, an additional artificial constraint. The results in French on the test set with diacritics are below English but within range overall. They are also below the results on the faulty test set in terms of SARI and compliance, but these differences may be due to the non-faulty test set being twice the size of the original one.

Overall, these results tend to indicate strong potential of multilingual LLMs for the SPRP task, factoring the lack of, or specific issues with, existing datasets in languages other than English.