

# Automating Easy Read Text Segmentation

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

{jcalleja, tetchegoyhen, adponce}@vicomtech.org

## Abstract

Easy Read text is one of the main forms of access to information for people with reading difficulties. One of the key characteristics of this type of text is the requirement to split sentences into smaller grammatical segments, to facilitate reading. Automated segmentation methods could foster the creation of Easy Read content, but their viability has yet to be addressed. In this work, we study novel methods for the task, leveraging masked and generative language models, along with constituent parsing. We conduct comprehensive automatic and human evaluations in three languages, analysing the strengths and weaknesses of the proposed alternatives, under scarce resource limitations. Our results highlight the viability of automated Easy Read text segmentation and remaining deficiencies compared to expert-driven human segmentation.

## 1 Introduction

Being able to access Easy Read (ER) text, also known as Easy-to-read, is critical for large segments of the population, including people with cognitive disabilities or suffering from learning difficulties, among others.<sup>1</sup> This type of text has two main characteristics. First, its content needs to consist of short sentences, using simple vocabulary and grammar as in text simplification, while also including an explanation of complex concepts in simpler terms as needed. Additionally, sentences need to be split into separate lines, following natural linguistic

boundaries, to facilitate reading.<sup>2</sup>

Although several aspects of ER text automation have been studied (Gala and Wilkens, 2020; Wilkens et al., 2022, 2024), to our knowledge automatic segmentation to meet ER readability requirements has not yet been explored in detail. The closest related domain where readability constraints require text segmentation along natural linguistic boundaries is subtitling. Despite this similarity, subtitling features specific constraints, such as the number of characters per line or characters per second, which are tied to the audiovisual environment. For ER text, these constraints do not apply and there is a clear gap in our knowledge regarding optimal methods for text segmentation in this context.

In this work, we explore several methods for automated segmentation and evaluate them on expert-crafted ER content in three languages, namely Basque, English and Spanish. Considering the scarcity of ER resources, we focus our study on methods able to generate segmentation hypotheses with either minimal or no training data. We design a scoring-based approach compatible with both pretrained constituency parsers and pretrained Masked Language Models (MLM). Additionally, we explore the use of pretrained generative Large Language Models (LLM), querying the models in zero-shot or few-shot fashion, as well as fine-tuning models on the limited available data.

We test the selected methods on datasets extracted from easy-to-read material collected from trustworthy sources, performing both automated and manual evaluations of the results. Our main results show that MLM-based scoring is a stable alternative across metrics, which may serve as a reasonable basis for ER text segmentation, although it still lags behind expert-crafted ER text.

Our main contributions can be summarised as

<sup>\*</sup>Equal contribution.

<sup>1</sup>The terms *Easy-to-read*/*Easy Read*, *Easy Language* and *Plain Language* are often used interchangeably, as they share a similar objective of facilitating the communication of information. There are however important differences between them, mainly that the target audience of ER content consists of people with reading difficulties, often with cognitive disabilities. ER text needs to be validated via focus groups, and the segmentation constraints are mainly meant to facilitate reading for this target group.

<sup>2</sup><https://www.inclusion-europe.eu/easy-to-read/>

follows:

- A first comprehensive assessment of automated segmentation for ER text generation.
- Novel approaches to ER text segmentation, including the use of generative LLMs under different modalities and a scoring method compatible with both constituency parsing and MLM modelling.
- Novel ER segmentation-centric datasets in Basque, English and Spanish, to support development in the field.<sup>3</sup>
- A comprehensive evaluation in terms of automated metrics and human qualitative assessments over system segmentation hypotheses as well as human references.

## 2 Related Work

Aspects of Easy Read text generation are related to work on text simplification, as it involves standard lexical and grammatical simplification processes (Saggion and Hirst, 2017; Alva-Manchego et al., 2020; Al-Thanyyan and Azmi, 2021). Our work is more specifically related to Easy Read content generation constraints, including sentence segmentation, and automated text segmentation approaches. We describe each aspect in turn below.

### 2.1 Easy Read

Easy Read is a text adaptation approach which aims at making information accessible to people with reading difficulties. A set of guidelines has been defined for ER document preparation across multiple languages,<sup>4</sup> covering text content, document layout, and validation of text comprehensibility by members of the target audience. The latter includes people with intellectual disabilities, prelingual hearing disabilities, aphasia, dyslexia or attention-deficit/hyperactivity disorder, among others. ER material can also be useful to anyone with reading difficulties, including people with limited knowledge of a language, for example.

In terms of text, the guidelines include recommendations at different levels. As for lexical choice, simple words should be used, and concepts deemed

<sup>3</sup>The datasets will be available under a CC-BY-NC-ND license at: <https://github.com/Vicomtech/EasyReadSeg>. As of this writing, note that efforts are still ongoing to secure sharing permissions for the content of all datasets.

<sup>4</sup>See, e.g., <https://www.inclusion-europe.eu/easy-to-read-standards-guidelines/>

too complex should be explained in simple terms. In terms of syntax, sentences should be simple and short, avoiding complex constructions. Finally, a sentence should fit into a single line wherever possible, but should otherwise be split into separate lines at natural linguistic boundaries, "where people would pause when reading out loud" as indicated in the Inclusion Europe English guidelines.<sup>5</sup>

In terms of length, a standard recommendation is to use sentences that comprise between 5 and 15 words, to enhance readability. Although this requirement is not specified in all ER guidelines,<sup>6</sup> it is generally considered standard practice in the field, for languages like English or Spanish at least.

These guidelines are meant to facilitate reading along several dimensions, with studies establishing empirical support for the use of easy-to-read texts, although some of the recommendations might deserve further investigation (Fajardo et al., 2014). González-Sordé and Matamala (2024) point out the relative lack of studies to firmly establish empirical support for ER guidelines.

### 2.2 Text Segmentation

To our knowledge, automated text segmentation has not been explored for ER text generation. As noted in the introduction, it is however an important aspect of subtitle generation, where subtitles should be segmented at naturally occurring linguistic boundaries to enhance readability, in addition to other subtitle-specific constraints. Subtitle segmentation has been shown to have an important impact on their readability (Perego et al., 2010; Rajendran et al., 2013).

Several approaches have been explored for automated subtitle segmentation. Thus, Álvarez et al. (2014) trained Support Vector Machine and Linear Regression models over professionally-created subtitles to predict subtitle breaks. This approach was later improved with the use of Conditional Random Fields (Álvarez et al., 2017). Other approaches jointly learn the generation of subtitle breaks within machine translation models (Matusov et al., 2019; Karakanta et al., 2020), an integrated approach which cannot be extended to the direct segmentation of text. Alternatively, Papi et al. (2022) proposed a multilingual segmenter which generates both text and breaks and may be trained on textual input only, or on joint text and audio data.

<sup>5</sup>*Op. cit.* footnote 4, pp. 17.

<sup>6</sup>This constraint is explicit in the Spanish ER norm.

Excepting the aforementioned integrated methods, most supervised approaches to subtitle segmentation could, in principle, translate to the ER use case, provided sufficient training data. Unfortunately, compared to subtitling, ER resources are scarce. To our knowledge, no ER datasets with segmented text are currently available, for any language. Furthermore, as shown in Section 4.1, the amount of ER data collectable from trustworthy sources is minimal to train machine learning models and, for most languages, simply non-existent. We will overcome these limitations by leveraging approaches to the task that rely on a limited number of samples, or unsupervised methods. Among the latter, we will notably adapt the approach of Ponce et al. (2023), which performs subtitle segmentation by computing the MLM-based likelihood of punctuation marks as an approximation to natural linguistic boundaries.

### 3 Methodology

The segmentation problem for ER text can be formulated as follows. Given a source sentence  $x$ , compute an ordered sequence of segments  $[s_1, s_2, \dots, s_n]$  such that the concatenation of all segments is identical to the source  $x$ , and all segments  $s_i$  follow natural linguistic boundaries, ending at a position in the sentence where a reader would naturally pause. Although intuitively clear, defining what *natural linguistic boundaries* or *natural pauses* actually mean can be a challenge. To avoid the need for a more specific definition, we will consider expert-crafted text segmentation as our reference standard and perform a manual evaluation of these reference datasets in Section 5.2.

We describe in turn below the different approaches we selected for ER segmentation. Further details regarding the specific models for each language and scenario are presented in Section 4.2.

#### 3.1 Scoring-based Segmentation

We first devised a simple approach compatible with any model that can assign a score to segmentation candidates between words in a given sentence. The sentence is tokenised into words, using white-space tokenisation. We then extract segmentation candidates within a fixed window, with minimum and maximum number of words per segment. The candidates are then sorted according to their score, as assigned by a given scoring function. In cases where there is insufficient text for segmentation,

the remaining text is selected as candidate. To avoid the pitfalls of selecting an early segmentation which may negatively affect the segmentation of the remainder of the input, we follow a beam-search approach and select the path with the best overall score, computed as the mean of all segments in the path. Appendix A provides further details on the segmentation algorithm.

The score of a given segment can be computed with any chosen method. In the next sections, we describe our two main approaches to compute segmentation scores, respectively based on constituency parsing and masked language modelling.

##### 3.1.1 Constituency Scoring

Our first scoring method is based on constituency parsing, using a pretrained model. Under this approach, segmentation scores are computed between any two words as the distance between the leaves associated to these words in the constituency tree. The intuition in this case is that, as the distance between nodes increases, the likelier it is that a break will preserve constituency groups. Additionally, given that the parsing tokenisation pipeline could either introduce new tokens or modify existing ones, and considering that our segmentation candidates should strictly align with white-space boundaries within the original text, we keep track of valid segmentation candidates throughout the scoring process, discarding any candidates that do not meet this criterion.

We chose constituency parsing, instead of alternatives such as chunking, as it provides higher parsing granularity and thus a higher number of segmentation points. One drawback of this approach is the limited coverage across languages. Nonetheless, it is still worth including this type of approach, as grammatical constituency is in line with what could be typically understood as delimiting natural linguistic boundaries.

##### 3.1.2 MLM Scoring

As our second scoring method, we adopted the unsupervised scoring approach to segmentation described in Ponce et al. (2023). We selected their best-performing variant, where segmentation candidates are assigned a score based on the MLM-predicted likelihood of a punctuation mark occurring in a mask inserted to the right of a given word.

The main difference between our approach and theirs lies in the segmentation algorithm itself. Whereas they follow a greedy approach, selecting

the best segmentation candidate within a segmentation window and proceeding recursively over the remainder of the text, we perform beam search to account for earlier segmentation choices that might not be optimal over the remaining text.

### 3.2 Generative LLM Segmentation

As our third approach, we queried pre-trained generative LLMs to generate segmented text. We used three different variants of this approach, described below. One notable potential drawback of this approach is the lack of guarantee that the generated text will not be altered beyond the introduction of line breaks. We measure this specific aspect in our experiments.

**Zero-shot.** We use simple prompts under a zero-shot approach with instruction-tuned LLMs and retrieve the model’s answers as our segmentation hypotheses. Given well-established LLM prompt sensitivity (Lu et al., 2022), we experimented with the following variants, similar in the three languages:<sup>7</sup>

- P1. Split the provided sentence into separate segments, inserting cuts where people would pause when reading out loud.
- P2. Split the provided sentence into separate segments that follow natural grammatical constituent limits.

Prompt P1 formulates the problem as per the guidelines referenced in Section 2.1, whereas prompt P2 refers to natural grammatical boundaries, avoiding references to reading pauses, which might be less directly interpretable by the model.

Both prompts are complemented with the following text, to include additional constraints regarding segment length and preservation of the original content: *"Each segment should have between 5 and 15 words. The content of the original sentence should be strictly maintained and no new information should be added"*.

**Few-shot** As a second variant, we performed few-shot prompting, providing the model with the same instructions as in our zero-shot approach and adding 5 examples of segmentation sampled from the development sets of the ER corpora, manually verified for correctness. As for zero-shot variants,

<sup>7</sup>See Appendix B for details on the prompts in Basque and Spanish.

one limitation of this approach is the lack of quality instruction-tuned models for most languages, to achieve proper segmentation.

**Fine-tuning** Finally, despite the relatively low amounts of training data, we fine-tuned base LLMs with LoRA (Hu et al., 2022), as this approach may overcome some of the limitations of zero-shot and few-shot approaches, such as prompt instability, i.e., result variation depending on prompt formulation. Further details on the setup in this scenario are provided in Appendix B.

## 4 Experimental Setup

In this section we describe the corpora we prepared for the task, the specific models selected for each approach and language, and our evaluation metrics.

### 4.1 Corpora

Our corpora were collected from three separate websites which feature expert-crafted ER text that adhere to the guidelines in the field. For English, we used ER articles from Inclusion Europe.<sup>8</sup> For Basque, we used articles from Irekia,<sup>9</sup> the transparency portal of the Basque Government, where a subset of the news are adapted for ER. Finally, for Spanish, we used both the Irekia website and articles from Plena Inclusión.<sup>10</sup>

We used in-house tools to scrape the data from the websites and perform boilerplate cleanup, notably removing enumerations presented via bullet points. We employed simple heuristics to identify segmentation, relying either on HTML elements whenever they permitted the identification of a line break, or sequences of uncased words in consecutive lines within text fields. All identified line breaks were then replaced with a segmentation marker and the text was split into separate sentences with scripts from the Moses toolkit (Koehn et al., 2007). Table 1 describes the core corpora statistics, in terms of number of sentences or segments, average sentence or segment length, and number or percentage of line breaks.

In all three languages, the majority of the sentences are segmented, still with notable differences between the most segmented (English) and the least segmented (Basque). The exact reasons for these differences are difficult to determine, as they could

<sup>8</sup><https://www.inclusion-europe.eu/category/etr/>

<sup>9</sup><https://www.irekia.euskadi.eus/lf/eu>

<sup>10</sup><https://www.plenainclusion.org/publicaciones/>

		EN	ES	EU
All	# sentences	4,441	14,353	10,956
	# breaks	3,840	11,904	7,832
	% breaks	86.47	82.94	71.49
no seg	# sentences	1,920	6,960	5,649
	Avg. sent. len.	7.62	7.73	6.11
w/ seg	# sentences	2,521	7,393	5,307
	Avg. # breaks	1.52	1.61	1.48
	Avg. sent. len.	14.98	19.19	14.62
	Avg. seg. len.	9.83	11.92	9.91

Table 1: Corpora statistics.

	EN	ES	EU
train	358	3,974	2,249
dev	500	500	500
test	1,500	1,500	1,500

Table 2: Data partition statistics (number of sentences)

be attributed to a difference in ER text adaptation style. One possible additional explanation is that in Basque, being an agglutinative language, there are simply fewer separate words in a sentence and therefore, fewer segmentation options.

Sentences without segmentation are typically short across the board, which is expected since splitting those sentences would result in separate lines that are too short. Sentences with segmentation abide by the general guidelines overall, with less than 15 words per sentence and segments under 10 words, on average. Interestingly, although some guidelines recommend a maximum of two lines per sentence, the actual data indicate an average of over 1.5 breaks per sentence, i.e. a preponderance of sentences split into more than two lines.

To train and evaluate the different segmentation approaches, we selected the subset of sentences with at least one segmentation break and partitioned the data as shown in Table 2. Note that that we performed additional filtering after identifying remnant noise in the data, mainly improperly split sentences with text following a final punctuation. Our main criterion for the partition was to reserve as large a set as possible for testing, instead of opting for slightly more training data with less accurate testing.

## 4.2 Models

**Constituency Parsers.** We selected Stanza (Qi et al., 2020) as our main tool for the task, using default pipelines and models for English and Spanish. Since there is no default coverage for Basque in Stanza, we integrated the Benepar EU model (Kitaev and Klein, 2018; Kitaev et al., 2019), which is trained on a filtered and adapted version of the Basque Constituency Treebank (Seddah et al., 2013). We adapted the Stanza code to incorporate this model, using the Stanza preprocessing and tagging pipeline available for Basque.

**Masked Language Models.** For English and Spanish, MLM scoring was computed with the multilingual BERT (Devlin et al., 2019) model;<sup>11</sup> for Basque, we used the IXAmBERT model (Otegi et al., 2020). We used default parameters and the *fill-mask* pipeline from the *transformers* library.<sup>12</sup>

**Generative Language Models.** For our zero-shot and few-shot experiments, we selected two different models. For all three languages, we used the Llama2-7B instruction model (Touvron et al., 2023).<sup>13</sup> To test a larger language model, we also used GPT-4 via the OpenAI API,<sup>14</sup> specifically model *gpt-4-0125-preview*. Considering the lack of public information regarding their training, along with potential test data leakage issues (Balocco et al., 2024), results with these models should be considered with all due caveats. For our fine-tuning experiments with LoRA, we used Llama2-7B in English and Spanish. For Basque, we selected the Latxa-7B model v1, a Llama2 model tuned on Basque corpora (Etxaniz et al., 2024).<sup>15</sup>

## 4.3 Evaluation Metrics

Standard segmentation metrics such as F1 accuracy assume that the hypotheses and references are identical, except for segmentation markers. However, this constraint would not necessarily hold for generative LLMs, since they may alter the original text when generating segmentation hypotheses. To address a similar issue with models that may generate imperfect text in subtitling, along with segmentation hypotheses, Karakanta et al. (2022)

<sup>11</sup><https://huggingface.co/google-bert/bert-base-multilingual-uncased>

<sup>12</sup><https://pypi.org/project/transformers/>

<sup>13</sup><https://huggingface.co/meta-llama/Llama-2-7b-chat-hf>

<sup>14</sup><https://pypi.org/project/openai/>

<sup>15</sup><https://huggingface.co/HitZ/latxa-7b-v1>

proposed the Sigma metric, which computes the ratio of achieved BLEU (Papineni et al., 2002) over an approximated upper-bound BLEU score. We include results with this metric to compare all models, along with BLEU scores to measure the impact of imperfect text generation. All metrics results were computed with the EvalSubtitle tool.<sup>16</sup>

#### 4.4 Human Evaluation

For our manual evaluations, we randomly sampled two subsets of 100 segmented sentences each: one sampled from the hypotheses of our overall best performing system, namely MLM; the other was sampled from the human segmentation references in the respective test sets for each language. Although human segmentation was produced by experts in the ER field, different segmentations are possible in general for a given sentence, and determining an optimal one can be subjective. Additionally, preliminary examination of the human references indicated cases of segmentation that did not adhere to the guidelines. We thus also aimed to measure the quality and consistency of the references in this evaluation.

To evaluate system hypotheses, there were 3 evaluators per language; for human segmentation, there were 6, 4 and 3 evaluators for English, Spanish and Basque, respectively. All evaluators were unpaid volunteers, proficient in the corresponding language. The evaluators were not ER experts, but ER segmentation guidelines are relatively straightforward, and we estimated that the concepts involved, such as grammatical boundaries, could be assessed by untrained speakers of a given language.

For each sentence, the evaluators had to answer 5 questions regarding the quality of the segmentation:

1. Do sentence splits occur where people would naturally pause? (disagree/mostly disagree/mostly agree/agree)
2. Do sentence splits occur along natural grammatical boundaries? (disagree/mostly disagree/mostly agree/agree)
3. Would you have split the sentence differently to facilitate reading? (yes/no)
4. Would you have used more splits to facilitate reading? (yes/no)
5. Would you have used fewer splits to facilitate reading? (yes/no)

More details on the evaluation protocol are provided in Appendix D.

<sup>16</sup><https://github.com/fyvo/EvalSubtitle>

## 5 Results

### 5.1 Automatic Evaluation

The comparative results on the Sigma, F1 and BLEU (excluding breaks) metrics, for all languages and models, are shown in table 3.

Taking all metrics into consideration, the best performing system overall was MLM-based segmentation, which achieved markedly better F1 scores in all languages, was competitive in Sigma, and achieved perfect BLEU scores, as expected for an approach that does not modify the input text. The Stanza variant was outperformed by MLM, with close Sigma results but significantly worse F1 scores, though still markedly better than all generative approaches. Stanza scores were lower in Basque overall, highlighting the dependency of this approach on the robustness of the available parsing tools for a given language.

For all generative approaches, F1 scores were drastically lower than those achieved by the scoring-based methods, a consequence, in part at least, of their potential text-altering characteristics. As indicated by BLEU scores, input text modification could reach significant amounts, with scores around or below 80 for Llama2-FS in ES and EU. The most preserving approach in this respect were fine-tuning of Llama2/Latxa and GPT4-FS, in all languages, with scores close to or above 99, although even small degrees of input text alteration would still imply a revision and eventual post-editing of the output text, in actual practice.

In terms of Sigma, the highest scores were achieved by GPT4 variants and fine-tuned Llama2/Latxa models. Results with the latter indicate that even relatively small amounts of training data can result in Sigma scores comparable to those of models of a much larger size such as the GPT4 models (purportedly).

The results were quite similar across language overall, with no notable deviation on metrics results per model variant. This was slightly unexpected, since all three languages are markedly different in terms of morphosyntax. It might be the case that even coarse-grained grammatical boundary identification is sufficient for the task, with most models capturing the core cases across languages.

Prompt variation did not have a significant impact in these experiments. This is also slightly surprising considering that one prompt refers to reading pauses and the other to grammatical boundaries. The common reference to splitting in the

	EN			ES			EU		
	Sigma	BLEU	F1	Sigma	BLEU	F1	Sigma	BLEU	F1
MLM	82.88	<b>100.00</b>	<b>40.80</b>	85.57	<b>100.00</b>	<b>42.16</b>	83.57	<b>100.00</b>	<b>41.71</b>
Stanza	80.35	<b>100.00</b>	32.44	83.79	<b>100.00</b>	36.32	71.27	<b>100.00</b>	27.19
Llama2-ZS-P1	<i>50.50</i>	<i>48.83</i>	11.44	71.54	<i>24.55</i>	<i>5.11</i>	<i>66.43</i>	<i>32.41</i>	1.89
Llama2-ZS-P2	56.32	57.40	13.73	63.33	<i>21.61</i>	7.00	70.66	73.47	<i>1.26</i>
GPT4-ZS-P1	<b>87.95</b>	77.08	<i>5.46</i>	86.88	93.27	7.20	85.52	88.72	6.32
GPT4-ZS-P2	84.72	86.04	8.68	87.23	95.62	8.68	<b>86.14</b>	81.97	7.53
Llama2-FS-P1	80.05	93.82	7.35	84.56	77.35	5.64	80.71	73.25	6.28
Llama2-FS-P2	79.85	94.11	7.11	84.36	78.04	6.62	80.40	80.97	4.89
GPT4-FS-P1	85.66	98.28	5.53	89.48	99.50	6.58	84.34	99.35	18.04
GPT4-FS-P2	85.84	97.99	6.37	<b>89.89</b>	99.51	6.45	83.20	99.04	10.42
Llama2/Latxa-FT-P1	85.06	99.80	14.63	89.37	99.50	7.55	86.02	99.78	12.63
Llama2/Latxa-FT-P2	84.60	99.73	18.32	89.86	99.68	6.64	86.10	99.58	15.17

Table 3: Comparative results on the ER segmentation test sets. BLEU scores were computed over text without breaks to measure input text alteration. Fine-tuned models were all based on Llama2, the Latxa variant being used for Basque. The models are grouped according to their main characteristics: scoring-based (MLM and Stanza), zero-shot (ZS), few-shot (FS) and fine-tuned (FT). P1 and P2 indicate the two prompt variants described in Section 3.2. Best results per metric and language are shown in bold, worst results in italics.

prompts might be the determining factor in comparable model behaviour, although a more detailed analysis would be needed to clarify this matter.

Finally, opting for zero-shot or few-shot did not markedly impact the GPT4 model on Sigma and F1, but the former led to the lowest results, by a large margin and across the board, for Llama2-7B. A zero-shot approach might be suboptimal for this type of smaller model on this task.

## 5.2 Human Evaluation

For the human evaluation, we selected the optimal segmentation system from the results of the previous section namely the scoring-based MLM variant. The results are shown in Table 4 and segmentation examples are provided in Appendix E.

As a first notable result, automated segmentation achieved lower scores than human segmentation overall for English and Spanish, outperformed by at least 10 percentage points in most cases. The system hypotheses were still valued positively overall, at 65% on grammatical boundaries and around 45% viewed as needing alternate segmentation. Note that even human reference segmentation was viewed as possibly requiring different segmentation in at least 33% of the cases, which indicates the difficult and partly subjective nature of the task.

For Basque, surprisingly, the human reference breaks were viewed as less respectful of reading

pauses, although still more grammatical overall. Alternate segmentation was also suggested in more cases for the human references. This highlights the difficulties of the task for human ER content creators as well, as text segmentation can be perceived as suboptimal overall due to individual reading preferences.

Both automatic and human segmentation received harsher scores in terms of reading pauses than grammatical boundaries. This could be due to the former being more subjective or dependent on individual readers than the latter. Inter-rater agreement was not markedly different though, with moderate agreement overall in both cases. This might thus simply indicate that the concept of natural reading pauses was more difficult to translate into proper segmentation.

Inter-rater agreement was higher on MLM segmentation overall, except for Basque where it was on a par or slightly higher than agreement on human references. This would tend to show that it was relatively easier to discriminate between correct and incorrect automated segmentation, with more salient errors when the automated system failed to generate a correct segmentation.

Finally, under-splitting was significantly more prevalent than over-splitting for both human and automated segmentation. This effect might be par-

MLM Segmentation			
	EN	ES	EU
Respects reading pauses?	59.33% / 1.93 [0.59]	59.67% / 2.08 [0.42]	57.33% / 2.18 [0.24]
Respects grammatical boundaries?	65.00% / 2.14 [0.56]	68.00% / 2.31 [0.50]	63.00% / 2.31 [0.28]
Alternate segmentation?	55.00% [0.56]	54.00% [0.51]	44.67% [0.44]
Use more breaks?	20.33% [0.53]	23.00% [0.46]	17.00% [0.29]
Use fewer breaks?	3.33% [0.69]	2.00% [0.32]	7.67% [0.25]
Human Segmentation			
	EN	ES	EU
Respects reading pauses?	70.20% / 2.36 [0.27]	69.75% / 2.49 [0.12]	51.00% / 2.10 [0.23]
Respects grammatical boundaries?	80.20% / 2.55 [0.33]	78.50% / 2.68 [0.06]	72.00% / 2.47 [0.20]
Alternate segmentation?	33.20% [0.32]	34.25% [0.29]	50.00% [0.39]
Use more breaks?	9.40% [0.21]	18.75% [0.21]	24.00% [0.13]
Use fewer breaks?	4.40% [0.26]	6.50% [0.07]	6.00% [0.12]

Table 4: Human evaluation results in English, Spanish and Basque, over segmentation samples from the MLM system and human experts. Answers to the first two questions are the percentage of *totally agree* answers and the average score on a [0-3] scale. For the other questions, we report the percentage of Yes. Krippendorff’s alpha inter-rater agreement is indicated between brackets.

tially due to the evaluators not suffering from reading disabilities, thus differing from what focus ER groups would value as optimal to facilitate reading. We leave this type of analysis for future work.

## 6 Error Analysis

**Segmentation Window.** To assess the impact of modelling the guidelines recommendations as a fixed window, we performed a grid search by varying the min and max parameters of the scoring-based models. Overall, the selected 5-15 window (F1: 40.80) was optimal in English, closely followed by the 1-10 window at 40.14. For Spanish, 1-10 achieved 45.22 and 5-10 yielded 44.09, both outperforming our default (42.16). For Basque, 5-15 at 41.71 was outperformed by 5-10 at 43.89. The complete results can be found in Appendix C.

Although improvements could have been achieved with different windows, these would have been somewhat marginal. Nonetheless, segmentation references outside our selected window would result in systematic errors for the scoring-based approach. There were 0.29%, 0.19% and 0.03% cases of segments over 15 words, for EN, ES and EU, respectively, of minimal impact. There were however 26.34%, 14.49% and 24.86% cases of segments under 5 words, for EN, ES and EU, respectively. These cases contributed significantly to the errors of the scoring-based variants.

Considering the relatively lax minimum number of words per segment in the reference sets, a possible solution would be to use a predefined segmentation window as a soft constraint rather than a hard limit on the segmentation candidates.

**Ungrammatical Segmentation.** This type of error is the most damaging for the task, as it directly impacts text readability. Although it would be difficult to uncover all failure patterns, we manually examined all cases in English where human evaluators assigned the worst score over both grammaticality and reading pause, in the human evaluation, which amounted to 36 examples.

Wrongly segmented complex named entities accounted for 25% of these errors, with cases such as: *Now, Thibeau works at Antwerp <seg> Management School as a researcher.*

Another pattern which also accounted for 25% of the cases were breaks after words that can act as a function word or as an independent word, e.g.: *One person working in the institution told her that <seg> she could live on her own..*

One additional pattern emerged, though only in 7% of the cases: split nominal coordination, influenced by the fact that a segment before a coordination can typically act in isolation, as in: *Policies can be a set of rules <seg> or guidelines to follow in or to achieve a specific goal.*

NER preprocessing, heuristics, or fine-tuning



the pretrained model on segmentation data could be alternatives worth exploring for these types of errors, at the cost of a more complex pipeline.

## 7 Conclusion

We described a comprehensive evaluation of automated segmentation for Easy Read text generation, a critical component of accessible content for people with reading difficulties.

We explored a wide range of methods, including scoring-based segmentation with masked language models or constituency parsing, and generative language models via zero-shot or few-shot prompting, and fine-tuning. Although limited ER resources are available overall, we collected data that supported minimal training and relatively large test sets in Basque, English and Spanish.

In terms of metrics, the unsupervised MLM scoring approach proved optimal overall, although generative models could achieved higher Sigma scores. The main deficiency of the latter was the generation of input altering text and low F1 scores for all variants. Fine-tuned LLMs generated the least amount of imperfect text, among generative variants, and might be a worthy alternative for the task.

Our human evaluations targeted both MLM-based and human segmentation, with the former lagging behind the latter in most cases, although the results were respectable overall across languages. We also identified segmentation error patterns from the MLM approach, which may provide research paths for future improvements.

## 8 Limitations

The evaluators that took part in our evaluation were not experts in the ER field, i.e. not trained to assess the readability of a given segmented sentence. Mitigating this limitation is the fact that the segmentation guidelines are relatively straightforward and the concepts involved, such as grammatical boundaries, can be assessed by untrained speakers of a given language. Nonetheless, ER text is specifically meant for people with reading difficulties, and an assessment of ER segmentation would gain from being performed by both experts in the field and focus groups. We leave this type of experimentation, which would involve more complex experimental protocols, for future studies.

## 9 Ethical Considerations

Automated methods for text segmentation are meant to facilitate reading for people with reading difficulties, including people with cognitive disabilities. These methods are experimental, with no guarantee that the output they produce will be error free or meet Easy Read standards. All automatically generated output should undergo human revision and be corrected as needed.

## Acknowledgments

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

## References

- Suha S Al-Thanyyan and Aqil M Azmi. 2021. Automated text simplification: a survey. *ACM Computing Surveys (CSUR)*, 54(2):1–36.
- Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. 2020. [Data-driven sentence simplification: Survey and benchmark](#). *Computational Linguistics*, 46(1):135–187.
- Aitor Álvarez, Haritz Arzelus, and Thierry Etchegoyhen. 2014. [Towards customized automatic segmentation of subtitles](#). In *Advances in Speech and Language Technologies for Iberian Languages*, pages 229–238, Cham. Springer International Publishing.
- Simone Balloccu, Patrícia Schmidtová, Mateusz Lango, and Ondrej Dusek. 2024. [Leak, cheat, repeat: Data contamination and evaluation malpractices in closed-source LLMs](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 67–93, St. Julian’s, Malta. Association for Computational Linguistics.
- 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.
- Julen Etxaniz, Oscar Sainz, Naiara Miguel, Itziar Aldabe, German Rigau, Eneko Agirre, Aitor Ormazabal, Mikel Artetxe, and Aitor Soroa. 2024. [Latxa: An open language model and evaluation suite for Basque](#).

- In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14952–14972, Bangkok, Thailand. Association for Computational Linguistics.
- Inmaculada Fajardo, Vicenta Ávila, Antonio Ferrer, Gema Tavares, Marcos Gómez, and Ana Hernández. 2014. Easy-to-read texts for students with intellectual disability: linguistic factors affecting comprehension. *Journal of applied research in intellectual disabilities*, 27(3):212–225.
- Núria Gala and Rodrigo Wilkens, editors. 2020. *Proceedings of the 1st Workshop on Tools and Resources to Empower People with READING Difficulties (READI)*. European Language Resources Association, Marseille, France.
- Mariona González-Sordé and Anna Matamala. 2024. Empirical evaluation of easy language recommendations: a systematic literature review from journal research in catalan, english, and spanish. *Universal Access in the Information Society*, 23(3):1369–1387.
- 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*.
- Alina Karakanta, François Buet, Mauro Cettolo, and François Yvon. 2022. **Evaluating subtitle segmentation for end-to-end generation systems**. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3069–3078, Marseille, France. European Language Resources Association.
- Alina Karakanta, Matteo Negri, and Marco Turchi. 2020. **Is 42 the answer to everything in subtitling-oriented speech translation?** In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 209–219, Online. Association for Computational Linguistics.
- Nikita Kitaev, Steven Cao, and Dan Klein. 2019. **Multilingual constituency parsing with self-attention and pre-training**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3499–3505, Florence, Italy. Association for Computational Linguistics.
- Nikita Kitaev and Dan Klein. 2018. **Constituency parsing with a self-attentive encoder**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. **Moses: Open source toolkit for statistical machine translation**. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. **Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland. Association for Computational Linguistics.
- Evgeny Matusov, Patrick Wilken, and Yota Georgakopoulou. 2019. **Customizing neural machine translation for subtitling**. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 82–93, Florence, Italy. Association for Computational Linguistics.
- Arantxa Otegi, Aitor Agirre, Jon Ander Campos, Aitor Soroa, and Eneko Agirre. 2020. Conversational question answering in low resource scenarios: A dataset and case study for basque. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 436–442.
- Sara Papi, Alina Karakanta, Matteo Negri, and Marco Turchi. 2022. **Dodging the data bottleneck: Automatic subtitling with automatically segmented ST corpora**. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 480–487, Online only. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing 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.
- Elisa Perego, Fabio Del Missier, Marco Porta, and Mauro Mosconi. 2010. **The cognitive effectiveness of subtitle processing**. *Media Psychology*, 13(3):243–272.
- David Ponce, Thierry Etchegoyhen, and Victor Ruiz. 2023. **Unsupervised subtitle segmentation with masked language models**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 771–781, Toronto, Canada. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. **Stanza: A python natural language processing toolkit for many human languages**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.

- Dhevi J. Rajendran, Andrew T. Duchowski, Pilar Orero, Juan Martínez, and Pablo Romero-Fresco. 2013. [Effects of text chunking on subtitling: A quantitative and qualitative examination](#). *Perspectives*, 21(1):5–21.
- Horacio Saggion and Graeme Hirst. 2017. *Automatic text simplification*, volume 32. Springer.
- Djamé Seddah, Reut Tsarfaty, Sandra Kübler, Marie Candito, Jinho D. Choi, Richárd Farkas, Jennifer Foster, Iakes Goenaga, Koldo Gojenola Gallettebeitia, Yoav Goldberg, Spence Green, Nizar Habash, Marco Kuhlmann, Wolfgang Maier, Joakim Nivre, Adam Przepiórkowski, Ryan Roth, Wolfgang Seeker, Yannick Versley, Veronika Vincze, Marcin Woliński, Alina Wróblewska, and Eric Villemonte de la Clergerie. 2013. [Overview of the SPMRL 2013 shared task: A cross-framework evaluation of parsing morphologically rich languages](#). In *Proceedings of the Fourth Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 146–182, Seattle, Washington, USA. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Rodrigo Wilkens, David Alfter, Rémi Cardon, and Núria Gala, editors. 2022. *Proceedings of the 2nd Workshop on Tools and Resources to Empower People with READING Difficulties (READI) within the 13th Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France.
- Rodrigo Wilkens, Rémi Cardon, Amalia Todirascu, and Núria Gala, editors. 2024. *Proceedings of the 3rd Workshop on Tools and Resources for People with READING Difficulties (READI) @ LREC-COLING 2024*. ELRA and ICCL, Torino, Italia.
- Aitor Álvarez, Carlos-D. Martínez-Hinarejos, Haritz Arzelus, Marina Balenciaga, and Arantza del Pozo. 2017. [Improving the automatic segmentation of subtitles through conditional random field](#). *Speech Communication*, 88:83–95.

## A Scoring Algorithm

The scoring algorithm described in Section A is presented in more detail in Algorithm 1.

Each sentence is tokenised in to an array of words, via white-space tokenisation, and passed to a scoring function, either one of the methods described in Section 3.1 in our experiments. The scoring function performs its own pre- and post-processing of the array of words, and returns an array of *scores* matching the number of words.

The *beam\_search* function receives the initial lists of words and segmentation scores. Both lists are aligned to have their indexes match, by including a score after the final word. The function applies a standard beam search to keep the number of candidates to the maximum indicated as *beam\_width*, effectively reducing the search space to the top *k* best scoring candidates.

The *get\_segmentation\_candidates* function, receives a list of words and its associated segmentation scores. The parameters *min\_words* and *max\_words* indicate predefined minimum and maximum number of words per segment, respectively. The function creates segmentation candidates within the window defined by these two parameters, where each candidate is assigned the score at its segmentation point, the words up to the segmentation point, and the remaining words and scores from this segmentation point on, if any. If the number of words to segment contains fewer than *min\_words*, the entire list of words is returned as a single segment and is assigned a penalty to decrease the likelihood of final segments with fewer than the required number of words.

## B Training and Inference

Fine-tuning was performed with LoRA (Hu et al., 2022), loading the base model in 8 bits, with the following parameters: *r*=8, *alpha*=16, *dropout*=0.0; targeted layers were Query and Value.

Other training parameters were as follows: *warmup\_steps*=100; *max\_steps*=10000; *optimizer*='adamw\_torch'; *batch\_size*=4; *gradient\_accumulation\_steps*=17; *learning\_rate*=3e-4; *max\_grad\_norm*=0.3; *weight\_decay*=0.01.

For generation, we used default parameters, namely: *max\_new\_tokens*=256, *temperature*=0.1, *top\_p*=0.9.

Each training sample was provided in the following format, where  $\alpha$  is the full sentence and  $\beta$  its corresponding segmented sentence:

INSTRUCTION: Split the provided sentence into separate segments, inserting cuts where people would pause when reading out loud. Each segment should have between 5 and 15 words. The content of the original sentence should be strictly maintained and no new information should be added.

INPUT:  $\alpha$

RESPONSE:  $\beta$

The following text was prepended to each sample prompt, in the corresponding language:

- *English*: Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request.
- *Spanish*: A continuación se muestra una instrucción que describe una tarea, junto con una entrada que proporciona más contexto. Escriba una respuesta que complete adecuadamente la solicitud.
- *Basque*: Jarraian, ataza bat deskribatzen duen jarraibide bat agertzen da, testuinguru gehiago ematen duen sarrera batekin batera. Idatzi eskaera behar bezala osatuko duen erantzun bat.

The following prompts were used for each language:

- Prompt P1:
  - *English*: Split the provided sentence into separate segments, inserting cuts where people would pause when reading out loud. Each segment should have between 5 and 15 words. The content of the original sentence should be strictly maintained and no new information should be added.

---

**Algorithm 1** Scoring-based Segmentation

---

```
function get_segmentation_candidates(words, scores)
  if length(words) < min_words then
    return {words : words, score : penalty, remaining_words :  $\emptyset$ , remaining_scores :  $\emptyset$ }
  end if
  for each index in range(length(words)) do
    if min_words > index < max_words then
      c  $\leftarrow$  new_candidate()
      c.score = scores[index]
      c.words = words[: index]
      c.remaining_words = words[index + 1 :]
      c.remaining_scores = scores[index + 1 :]
      candidates.append(c)
    end if
  end for
  return candidates
end function

function beam_search(words, scores, beam_width)
  candidates  $\leftarrow$  get_segmentation_candidates(words, scores)
  while true do
    new_candidates  $\leftarrow$  [ ]
    for each c in candidates do
      sc  $\leftarrow$  get_segmentation_candidates(c.remaining_words, c.remaining_scores)
      new_candidates.append(sc)
    end for
    sort_by_score(new_candidates)
    candidates  $\leftarrow$  new_candidates[: beam_width]
    best  $\leftarrow$  candidates[0]
    if best.remaining_words =  $\emptyset$  then
      return best
    end if
  end while
end function
```

---

- *Spanish*: Parte la frase proporcionada en diferentes segmentos, insertando cortes donde las personas se pararían al leerla. Cada segmento debería tener entre 5 y 15 palabras. El contenido de la frase original debería mantenerse estrictamente y no se debería añadir información nueva.
  - *Basque*: Zatitu emandako esaldia segmentu desberdinetan, jendea irakurtzean geldituko litzatekeen lekuetan moztuz. Segmentu bakoitzak 5 eta 15 arteko hitz kopurua izan beharko luke. Esaldiaren edukia zorrozki mantendu beharko litzateke eta ez litzateke informazio berririk gehitu beharko.
- Prompt P2:
    - *English*: Split the provided sentence into separate segments that follow natural grammatical boundaries. Each segment should have between 5 and 15 words. The content of the original sentence should be strictly maintained and no new information should be added.
    - *Spanish*: Parte la frase proporcionada en diferentes segmentos que respeten fronteras lingüísticas naturales. Cada segmento debería tener entre 5 y 15 palabras. El contenido de la frase original debería mantenerse estrictamente y no se debería añadir información nueva.
    - *Basque*: Emandako esaldia muga linguistiko naturalak errespetatzen dituen segmentuetan zatitu. Segmentu bakoitzak 5 eta 15 arteko hitz kopurua izan beharko luke. Esaldiaren edukia zorrozki mantendu beharko litzateke eta ez litzateke informazio berririk gehitu beharko.

Basque and English models were trained on 1 L40 with 48GB of RAM; Spanish models on 1 L40S with 48GB of RAM. Checkpoints were saved every 30 steps. The best checkpoints were achieved at steps 180 and 210 for prompts P1 and P2, respectively, for Basque; 330 and 390 for Spanish; and step 90 for both prompts in English.

For few-shot and zero-shot, we used the default Llama2 template, using the system role for the instruction and the user role for the sentence to be segmented. For few-shot prompting, 5 examples of segmentation were shown after the instruction, using the user and assistant roles for the input and output examples.

## C Grid Search

The complete results of the grid-search described and discussed in Section 6 are shown in Figure 1.




Figure 1: Grid search F1 scores for scoring-based methods MLM and Stanza. The x axis indicates minimum and maximum number of words of the segmentation window.

## D Human Evaluation Protocol

To evaluate system hypotheses, there were 3 evaluators per language; for human segmentation, there were 6, 4 and 3 evaluators for English, Spanish and Basque, respectively. All evaluators were unpaid volunteers, proficient in the corresponding language.

The samples and associated questions were provided within an ad-hoc Web-based environment developed for the task, shown in Figure 2. The following instructions were provided to the evaluators:

---

We call for measures that protect people with intellectual disabilities during and after the Coronavirus emergency.

1. *Do sentence splits occur where people would naturally pause?*  
 Completely disagree  Mostly disagree  Mostly agree  Completely agree
2. *Do sentence splits occur along natural grammatical boundaries?*  
 Completely disagree  Mostly disagree  Mostly agree  Completely agree
3. *Would you have split the sentence differently to facilitate reading?*  
 No  Yes
4. *Would you have used more splits to facilitate reading?*  
 No  Yes
5. *Would you have used fewer splits to facilitate reading?*  
 No  Yes

Previous

Next

Figure 2: Web interface for the qualitative evaluation

The general guidelines to facilitate reading via segmentation are as follows:

- Insert line breaks (splits) that follow natural linguistic boundaries.
- Try to ensure that each line has a minimum of 5 words and a maximum of 15. This recommendation is not mandatory and it is advised to prioritise what would make reading easier.

The meaning of each question in the evaluation is indicated below and guidelines are provided to respond appropriately:

1. Do the splits in the sentence occur where a person would pause when reading it?
  - This question aims to assess whether the splits correspond to positions in the sentence where it would be natural for a native speaker to pause while reading.
  - Example of a natural split as a pause:
    - People participated in the experiment with enthusiasm during the day.
  - Example of a less natural split as a pause:
    - People participated in the experiment with enthusiasm during the day.
  - In the event that there is more than one break in the sentence:
    - If all splits are natural, indicate: Totally agree.
    - If the majority of splits are natural, indicate: Partially agree.
    - If the majority or half the splits are unnatural, indicate: Partially disagree.

- If none is natural, indicate: Totally disagree.
2. Do the splits in the sentence preserve linguistic boundaries?
- This question aims to assess whether the splits preserve grammatical boundaries.
  - Example of a natural split in terms of grammatical boundaries:
    - People participated  
in the experiment  
with enthusiasm  
during the day.
  - Example of a less natural split in terms of grammatical boundaries:
    - People  
participated in the  
experiment with  
enthusiasm during the day.
  - In the event that there is more than one split in the sentence:
    - If all splits preserve grammatical boundaries, indicate: Totally agree.
    - If the majority of splits preserve grammatical boundaries, indicate: Partially agree.
    - If the majority or half the splits do not preserve grammatical boundaries, indicate: Partially disagree.
    - If none preserve grammatical boundaries, indicate: Totally disagree.
3. Would you have used different splits in the sentence to make it easier to read?
- Answer YES if you would have modified the splits in any way:
    - Moving the current splits to another position and/or
    - Adding splits and/or
    - Removing splits.
  - Answer NO if you would not make any change to the current splits.
4. Would you have used more splits in the sentence to make it easier to read?
- Answer YES if you had used more splits.
  - Answer NO otherwise.
5. Would you have used fewer splits in the sentence to make it easier to read?
- Answer YES if you had used fewer splits.
  - Answer NO otherwise.

## **E Segmentation Examples**

Tables 5, 6 and 7 provide examples of correct MLM-scoring segmentation, incorrect MLM-scoring segmentation, and incorrect human segmentation, respectively.



Language	Segmentation
EN	It means that people with disabilities could not have access to healthcare.
	In order to find a job, we also need to access vocational training.
	To this date, the NGO has about 60 volunteers, who help about 60 children with disabilities
ES	También se han debatido las enmiendas que han realizado los demás partidos políticos. <i>The amendments that were made by the other political parties were also debated there.</i>
	En los cursos que se hacen en la Escuela de Pastores la cuarta parte de los alumnos han sido mujeres y esa presencia es importante y hay que hacerla visible. <i>In the courses that are held at the Escuela de Pastores the fourth part of the students there have been women and that presence is important and has to be made visible.</i>
	La lectura fácil es útil para muchas personas con discapacidad intelectual o del desarrollo. <i>Easy reading is useful for many people with intellectual or developmental disabilities.</i>
EU	Horrela, pertsonak elkarrizketa eta akordioak egiten dituzte, bata besteen artean aurka egin ordez. <i>This way, people make conversations and agreements, instead of opposing each other.</i>
	Bilera honekin indartu nahi da bi erakundeen arteko lankidetzeta. <i>This meeting is intended to strengthen collaboration between the two organizations.</i>
	Sistema horrek herritarrei abisua ematen die larrialdi handi edo hondamendi kasuetan. <i>This system warns citizens in case of major emergency or disaster.</i>

Table 5: Examples of MLM scoring-based segmentation viewed as correct

Language	Segmentation
EN	We talked about deinstitutionalisation with fifteen countries of the European Union.
	People working in institutions did not have enough protections (masks, gloves).
	EPSA members talked about asking for more money to solve this problem.
ES	Un delito por ejemplo puede ser que te hayan pegado. <i>A crime, for example, could be that you have been hit.</i>
	Si necesitas más información, te recomendamos leer el reglamento original. <i>If you need more information, we recommend reading the original regulations.</i>
	Pero que es más importante aún, que los cocineros y cocineras del jurado se hayan comprometido con su proyecto Zero Foodprint. <i>But what is even more important is that the chefs on the jury have committed to their Zero Foodprint project.</i>
EU	Azkenik, sailburuak esan du istripuarekin lotutako prozesu administratibo eta judicial guztiak aurrera jarraitzen dutela. <i>Finally, the minister said that all administrative and judicial processes related to the accident proceed.</i>
	Dirua gehiagotu da Zinematografiaren eta Arteen Institutuarekin (ICCA) hitzarmen bat sinatuko delako. <i>The money has increased because an agreement will be signed with the Institute of Cinematography and Arts (ICCA).</i>
	Euskadik ez du ahaztuko izan zuten gizatasuna eta elkartasuna. <i>Euskadi will not forget their humanity and solidarity.</i>

Table 6: Examples of MLM scoring-based segmentation viewed as incorrect

Language	Segmentation
EN	Atempo is a company that helps people with disabilities.
	The conference was about the 10 years of deinstitutionalization, in Europe.
	Paul Alford did not have a choice on whom to live with.
ES	Conoce cómo funciona el Parlamento de La Rioja, su historia y curiosidades. <i>Learn how the Parliament of La Rioja works, its history and curiosities.</i>
	Eso quiere decir que es fácil de entender. <i>That means it is easy to understand.</i>
	Erkoreka ha dicho que el acuerdo es posible pero es necesario que el Gobierno español también avance y cumpla con los compromisos que han tomado. <i>Erkoreka has said that the agreement is possible but it is necessary for the Spanish Government to also move forward and fulfill the commitments they have made.</i>
EU	Sektore horiek ere lotuta daudelako klimarekin. <i>Because these sectors are also linked with the climate.</i>
	Bonoak erosten dituzten pertsonak 2021eko maiatzaren 31ra arte erabili ahal izango dituzte. <i>People who buy vouchers will be able to use them until May 31, 2021.</i>
	Guztira, miloi erditik gora euro emango dira, lau diru-laguntzatan banatuta. <i>In total, more than half a million euros will be given, divided into four grants.</i>

Table 7: Examples of human references viewed as incorrect