

BiSECT: Learning to Split and Rephrase Sentences with Bitexts

Joongwon Kim^{1*}, Mounica Maddela^{2*}, Reno Kriz,^{1,3} Wei Xu,² Chris Callison-Burch¹

¹Department of Computer and Information Science, University of Pennsylvania,

²School of Interactive Computing, Georgia Institute of Technology

³ Human Language Technology Center of Excellence, Johns Hopkins University

{jkim0118, ccb}@seas.upenn.edu, {mmadela3, wei.xu}@cc.gatech.edu, rkriz1@jh.edu

Abstract

An important task in NLP applications such as sentence simplification is the ability to take a long, complex sentence and split it into shorter sentences, rephrasing as necessary. We introduce a novel dataset and a new model for this ‘split and rephrase’ task. Our BiSECT training data consists of 1 million long English sentences paired with shorter, meaning-equivalent English sentences. We obtain these by extracting 1-2 sentence alignments in bilingual parallel corpora and then using machine translation to convert both sides of the corpus into the same language. BiSECT contains higher quality training examples than previous Split and Rephrase corpora, with sentence splits that require more significant modifications. We categorize examples in our corpus, and use these categories in a novel model that allows us to target specific regions of the input sentence to be split and edited. Moreover, we show that models trained on BiSECT can perform a wider variety of split operations and improve upon previous state-of-the-art approaches in automatic and human evaluations.¹

1 Introduction

Understanding long and complex sentences is challenging for both humans and NLP models. NLP tasks like machine translation (Pouget-Abadie et al., 2014; Koehn and Knowles, 2017) and dependency parsing (McDonald and Nivre, 2011) tend to perform poorly on long sentences. Text simplification (Zhu et al., 2010; Xu et al., 2015) is often formulated with a specific step to break longer sentences into shorter sentences. This task is referred to as Split and Rephrase (Narayan et al., 2017).

Several past efforts have created Split and Rephrase training sets, which consist of long, complex input sentences paired with multiple shorter

* Equal contribution.

¹Our code and data are available at <https://github.com/mounicam/BiSECT>.

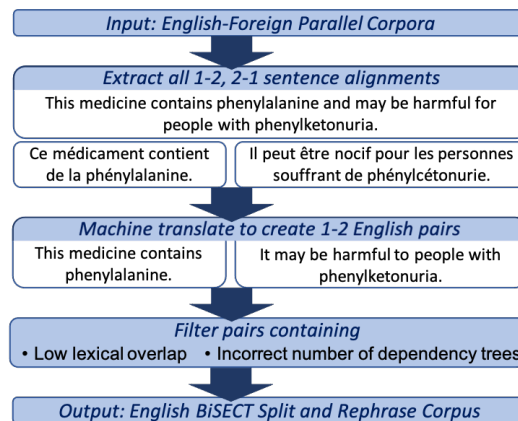


Figure 1: The process of creating the English BiSECT Split and Rephrase corpus.

sentences that preserve the meaning of the input sentence. Narayan et al. (2017) introduced the WEBSPLIT corpus based on decomposing a long sentence into RDF triples (a form of semantic representation), and generating shorter sentences from subsets of these triples. However, the reliance on RDF triples and a limited vocabulary results in unnatural expressions (Botha et al., 2018) and repeated syntactic patterns (Zhang et al., 2020a).

More recently, the WIKISPLIT corpus (Botha et al., 2018) was introduced. It contains one million training examples of sentence splitting that were mined from the revision history of English Wikipedia. While this yields an impressive number of training examples, the data are often quite noisy, with around 25% of WIKISPLIT pairs containing significant errors (detailed in §3.2). This is because Wikipedia editors are not only trying to split a sentence, but also often simultaneously modifying the sentence for other purposes, which results in changes of the initial meaning.

In this paper, we introduce a novel methodology for creating Split and Rephrase corpora via bilingual pivoting (Wieting and Gimpel, 2018; Hu et al., 2019b). Figure 1 demonstrates the process. First,

we extract all 1-2 and 2-1 sentence-level alignments (Gale and Church, 1993) from bilingual parallel corpora, where a single sentence in one language aligns to two sentences in the other language. We then machine translate the foreign sentences into English. The result is our BiSECT corpus.

Split and Rephrase corpora, including BiSECT, contain pairs with variable amounts of rephrasing. Some pairs only edit around the split location, while others require more involved changes to maintain fluency. In this work, we leverage this knowledge by introducing a classification task to predict the amount of rephrasing required, and a novel model that targets that amount of rephrasing.

The main contributions of this paper are:

- We introduce BiSECT, the largest multilingual Split and Rephrase corpus. BiSECT contains 938K English pairs, 494K French pairs, 290K Spanish pairs, and 186K German pairs.
- We show that BiSECT is higher quality than WIKISPLIT, that it contains a wider variety of splitting operations, and that models trained with our resource produce better output for the Split and Rephrase task.
- We introduce a novel classification task to identify the types of sentence splitting outputs based on how much rephrasing is necessary.
- We develop a novel Split and Rephrase model that accounts for these classifications to control the amount of rephrasing.

2 Related Work

The idea of splitting a sentence into multiple shorter sentences was initially considered a sub-task of text simplification (Zhu et al., 2010; Narayan and Gardent, 2014). However, the structural paraphrasing required to split a sentence makes for an interesting problem in itself, with many downstream NLP applications. Thus, Narayan et al. (2017) proposed the Split and Rephrase task, and introduced the WEBSPLIT corpus, created by aligning sentences in WebNLG (Gardent et al., 2017). WEBSPLIT contains duplicate instances and phrasal repetitions (Aharoni and Goldberg, 2018; Botha et al., 2018), and most splitting operations can be trivially classified (Zhang et al., 2020a), so subsequent Split and Rephrase corpora have been created to improve training (Botha et al., 2018) and evaluation (Sulem et al., 2018; Zhang et al., 2020a). The main work we compare against is WIKISPLIT, a corpus created by extracting split sentences from Wikipedia

edit histories (Botha et al., 2018). Concurrent work used a subset of WIKISPLIT to focus on sentence decomposition (Gao et al., 2021). While this approach is able to both extract many potential sentence splits and transfer across languages, edited sentences do not necessarily have to retain the same meaning. In contrast, our corpus BiSECT is created from aligned parallel documents.

Bilingual corpora is generally leveraged for monolingual tasks with bilingual pivoting (Barnard and Callison-Burch, 2005), which assumes that two English phrases that translate to the same foreign phrase have similar meaning. This technique was used to create the Paraphrase Database (Ganitkevitch et al., 2013; Pavlick et al., 2015), a collection of over 100 million paraphrase pairs, and to improve neural approaches for sentential paraphrasing (Mallinson et al., 2017; Wieting and Gimpel, 2018; Hu et al., 2019a,b) and sentence compression (Mallinson et al., 2018).

In introducing the Split and Rephrase task, Narayan et al. (2017) also reports the performance of several baseline models, where the strongest is an LSTM-based model. Subsequent works have improved performance using a copy-attention mechanism (Aharoni and Goldberg, 2018). We instead start with a BERT-initialized transformer model (Rothe et al., 2020), and train it with an adaptive loss function to emphasize split-based edits. Concurrent work also introduced an additional neural graph-approach for Split and Rephrase (Gao et al., 2021).

3 BiSECT Corpus

To address the need of Split and Rephrase data that is both meaning preserving and sufficient in size for training, we present the BiSECT corpus.

3.1 Corpus Creation Procedure

The construction of the BiSECT corpus relies on leveraging the sentence-level alignments from OPUS (Tiedemann and Nygaard, 2004), a publicly available collection of bilingual parallel corpora over many language pairs. While most of the translated sentences in OPUS are aligned 1-1, i.e., one sentence in Language A is mapped to one sentence in Language B , there are many aligned pairs consisting of multiple sentences from either A or B . This is a result of natural variation in the process of human translation. Sentence alignment algorithms (Gale and Church, 1993) match 1-1, 2-1, and 1-2

Dataset	Pivot Lang.	Domain	1-2 & 2-1 Alignments		Length	
			total (count/%)	after filtering	Long	Split
CCALIGNED	fr	web crawl	559,826 (20.9%)	203,780	36.0	20.1
EUROPARL	fr	European Parliament	153,220 (5.72%)	57,473	45.6	23.4
10 ⁹ FR-EN	fr	newswire	624,381 (23.31%)	264,203	41.8	22.5
PARACRAWL	fr,de,es,nl,it,pt	web crawl	1,212,982 (45.29%)	405,612	38.5	19.7
UN	fr,es,ar,ru	United Nations	113,840 (4.25%)	64,690	45.5	24.4
EMEA	fr	European Medicines Agency	5,719 (0.21%)	1,056	34.1	19.7
JRC-ACQUIS	fr,de	European Union	8,358 (0.31%)	6,237	51.8	26.6

Table 1: Datasets from OPUS that were used to create the English version of BiSECT. The training set consists of five corpora in the upper part of the table, while the two corpora in the lower part are used for the development and test sets. We also report the token length of the long sentence and that of the individual split sentences.

alignments in bitext. We extract all 1-2 and 2-1 sentence alignments from parallel corpora, where A is English and B is one of several foreign languages.

Next, the foreign sentences are translated into English using Google Translate’s Web API service² to obtain English sentence alignments between a single long sentence l and two corresponding split sentences $s = (s_1; s_2)$. As the alignment information provided by OPUS is based on the presence of a sentence-breaking punctuation, there are noisy alignments where l contains a pair of sentences instead of one complex sentence. These noisy alignments belong to two categories: two sentences pasted contiguously without any space around the sentence-breaking delimiter and two independent sentences joined by a space without any punctuation. For the first case, we remove l and its corresponding splits when it contains a token with a punctuation after the first two and before the last two alphabetic characters. For the second case, we generate a dependency tree³ for l and discard l if it contains more than one unconnected component.

Moreover, we remove the misalignment errors based on lexical and semantic overlap. We compute lexical overlap ratio r as follows:

$$r = \min\left(\frac{|\mathcal{L}_l \cap \mathcal{L}_{s_1}|}{|\mathcal{L}_{s_1}|}, \frac{|\mathcal{L}_l \cap \mathcal{L}_{s_2}|}{|\mathcal{L}_{s_2}|}, \frac{|\mathcal{L}_l \cap (\mathcal{L}_{s_1} \cup \mathcal{L}_{s_2})|}{|\mathcal{L}_{s_1} \cup \mathcal{L}_{s_2}|}\right);$$

where \mathcal{L}_l , \mathcal{L}_{s_1} and \mathcal{L}_{s_2} denote the sets of lemmatized tokens in l and $(s_1; s_2)$, respectively. We consider an aligned pair valid if $r \geq 0.25$ and l , s_1 and s_2 all contain a verb. We discard invalid pairs. We also remove $(l; s)$ pairs with length-penalized

BERTScore < 0.4 (Zhang et al., 2020b; Maddela et al., 2021).⁴

We repeat this process over all available parallel corpora for each English-Foreign language pair, resulting in 938,102 filtered English-English pairs. An important characteristic of BiSECT to note is that its size can be further increased with the addition of new parallel corpora on OPUS, processed in the method described above.

Table 1 breaks down the OPUS corpora and parallel languages used in creating the English version of BiSECT. For the testing set, a different set of corpora is used from the training set to prevent domain overlap. Moreover, the choice of corpus is based on the number of alignments extracted from each corpus. We choose corpora of relatively smaller sizes for development and testing to avoid a loss of size in the training set. To demonstrate our approach can be extended to other languages, we also create BiSECT corpora for French, Spanish, and German, using English as the pivot language. Corpus statistics of non-English languages are given in Appendix G.

3.2 Comparison to Existing Corpora

Corpus Statistics. Besides corpus size, we are interested in the amount of rephrasing (indicated by %new) and the syntactic complexity of sentences (approximated by length). In Table 2, we compare BiSECT with previous split and rephrase corpora, including WIKISPLIT (Botha et al., 2018), WEBSPLIT (Narayan et al., 2017; Aharoni and Goldberg, 2018), HSplit-Wiki (Sulem et al., 2018), Contract and Wiki-BM (Zhang et al., 2020a). BiSECT is comparable in size with WIKISPLIT, while impor-

⁴We also tried to fix the grammatical errors in the (l, s) pairs using GECToR (Omelianchuk et al., 2020). However, GECToR introduced minimal one word changes that did not help in improving the quality of the data.

²<https://pypi.org/project/googletrans/>

³We generate dependency trees using Spacy.

Corpus	#pairs	#unique	%new	Length	
				Long	Split
HSPLIT-WIKI [†]	1436	359	33.9	22.6	14.3
CONTRACT [†]	659	406	10.7	39.7	22.9
WIKI-BM [†]	720	403	8.9	29.2	16.1
WEBSPLITV1.0	1.06M	17k	32.1	34.3	30.2
WIKISPLIT	999K	999K	15.5	33.4	19.0
BISECT (this work)	938K	938K	34.6	40.1	20.6

Table 2: Comparison of Split and Rephrase corpora. We compute the number of aligned pairs (**#pairs**); number of unique long sentences l (**#unique**); the percentage of new words added to s compared to l (**%new**), and the average token **Length** of l and that of the individual split sentences. [†] marks crowdsourced corpora.

tantly containing longer aligned sentence pairs and a higher %new score, indicating that BIASECT contains more complex pairs with significantly more rephrasing (see also examples in Tables 3 and 4).

Manual Quality Assessment. While BIASECT does not suffer from meaning-altering edits like WIKISPLIT does, a potential concern is the error induced from translating a foreign text to English. Thus, we perform a manual assessment of corpus quality by comparing 100 randomly selected pairs from both BIASECT and WIKISPLIT corpora. We categorize each example ($l; s$) into two groups: (1) *high-quality pairs*, where both l and s are grammatical, l consists of exactly one sentence, and s contains exactly two sentences; and (2) *significant errors*, where the pair contains drastic errors impacting its usability. Table 3 shows the results of the manual inspection. When compared with WIKISPLIT, BIASECT contains significantly more high-quality pairs, while containing fewer pairs with significant errors. Pairs containing unsupported and deleted details are comparable across corpora, though WIKISPLIT skews more towards adding unsupported information, which is consistent with previous work (Zhang et al., 2020a).

Moreover, we take 100 random samples from the German BIASECT corpus and perform manual inspection. We chose German because translating to/from German is notoriously challenging for translation systems (Twain, 1880; Collins et al., 2005). As shown in Table 3, German BIASECT still contains 77% high-quality pairs.

3.3 Categorization for Split and Rephrase

One aspect of the Split and Rephrase task that has received little attention, outside of Zhang et al. (2020a), is the amount of rephrasing that occurs in

each instance, and more specifically the syntactic patterns involved in this rephrasing. Unlike more open-ended language generation tasks, the structural paraphrasing involved in Split and Rephrase is likely to be relatively consistent across domains, thus identifying these patterns is a critical step towards further improvement of neural-based approaches. In this work, we consider three major categories, and break down each of these further into more specific syntactic patterns. The categories are derived from the entire dataset, spanning the domains of web, newswire, medical and legal text, and others.

The first group involves *Direct Insertion*, when a long sentence l contains two independent clauses, and requires only minor changes in order to make a fluent and meaning-preserving split s . Within this category, we identify two sub-categories: *Colon/Semicolon*, which occurs when the clauses are connected by a colon or semicolon; and *Conjunction with subject*, where the clauses are connected by a conjunction, and the second clause contains an explicit subject. The second group involves *Changes near Split*, when l contains one independent and one dependent clause, but modifications are restricted to the region where l is split. Within this category, we identify four sub-categories: instances containing a *conjunction without subject*, which involves two clauses connected by a conjunction, but the second clause does not have an explicit subject; instances that contain a *gerund*, followed by an adjectival clause, adverbial clause, or prepositional phrase; instances that involve an explicit *subordinate clause*; and instance that contain a *concluding relative clause*. Finally, the third major group involves *Changes across Sentences*, where major changes are required throughout l in order to create a fluent split s . The main subcategory within this group involves a *preceding relative clause*, followed by a comma.

Table 4 presents the examples and prevalence of each category in WIKISPLIT and BIASECT, computed using a manual inspection of 100 random examples from each corpus. BIASECT contains significantly more instances that require changes across the sentence to form a high-quality split. To assess the relative difficulty of these categories, we analyze the quality of sentence splits generated by **DisSim** (Niklaus et al., 2019), a rule-based sentence splitter, on these 200 selected examples. DisSim splits the source sentence recursively using 35

Original Text	Split Text	WIKI	BISECT en	BISECT de
<i>High-Quality Split and Rephrase pairs</i>		73%	85%	77%
An additional advantage is that a shorter ramp can be used, thereby reducing weight and improving the rear view of the driver. (de→en)	Another advantage is that a shorter ramp can be used. This saves weight and improves the look of the rear of the vehicle.			<i>Perfect pairs</i>
Bitte geben Sie hier Ihre E-Mail-Adresse ein und wir senden Ihnen anschließend einen Link zu, mit dem Sie Ihr Passwort zurücksetzen können. (en→de)	Bitte geben Sie unten Ihre E-Mail-Adresse ein. Wir senden Ihnen einen Link per E-Mail, mit dem Sie ein neues Passwort erstellen können.	51%	63%	53%
Its many novel features ensure that it is easy to use correctly, making it suitable for all patients regardless of disease severity, in the elderly and for children. (de→en)	Its numerous control mechanisms ensure that the Novolizer is easy to use correctly. This makes it suitable for all patients regardless of the severity of the disease, for older patients and for children.			<i>Unsupported Details</i>
Every day, pedestrians take risks by working near mobile machinery and every day, accidents cost businesses dearly. (fr→en)	Every day, men take risks with machines . And every day accidents cost businesses dearly.			<i>Deleted Details</i>
<i>Pairs with significant errors</i>		27%	15%	23%
A little after the issue of Tosatti's book, Rizzoli published another volume on Fatima, this time a book-interview with Cardinal Bertone, edited by Vatican expert Giuseppe De Carli. (de→en)	Shortly after the publication of Tosatti's book, the Italian publisher Rizzoli published another book on Fatima. An interview book with Cardinal Bertone, edited by the Vaticanist Giuseppe De Carli.	10%	5%	12%
The children concoct many plans to lure Boo Radley out of his house for a few summers until Atticus make not true out , and they become "engaged." (WikiSplit)	The children concoct many plans to lure Boo Radley out of his house for a few summers until Atticus makes them stop. Dill promises to marry Scout , and they become "engaged."			<i>Multiple Errors</i>
Dann setzt unser Destillateurmeister die Brennblase in Gang und destilliert unter den Augen der Teilnehmer einen Berlin Dry Gin, der natürlich am Ende der Veranstaltung verkostet werden kann. (en→de)	Distiller legt den noch in Bewegung in Bewegung und Destillieren unter den Augen der Teilnehmer ein Berliner trockener Gin, der natürlich am Ende der Veranstaltung geschmeckt werden kann. Und während die noch Blasen, tauchen die Teilnehmer in die Welt des Gin ein.	17%	10%	11%

Table 3: Examples of high-quality and noisy sentence splits in the BISECT corpus. Some examples have **minor adequacy/fluency issues** (not uncommon in most existing monolingual parallel corpora) and are still usable, while a small portion (15%) contain more **significant errors**. Prevalence of each category is calculated based on 100 manually inspected pairs from WIKISPLIT (Botha et al., 2018) and English/German BISECT (our work).

hand-crafted rules based on a syntactic parse tree. DisSim produces disfluent sentence splits 34% of the time, and performs no splitting 9% of the time. For the *Changes near Split* and *Changes Across Sentence* categories, the number of erroneous splits increases to 55% and 63%, respectively. Although rules correctly identify the location of sentence splits, they fail to effectively modify sentences requiring more expansive rephrasing.

4 Our Model

The BISECT corpus contains a significant amount of paraphrasing along with sentence splitting, and models trained on BISECT tend to alter the lexical choices made in the input sentence. Although this is desirable in some situations, like for the task of sentence simplification, sometimes it can alter the meaning of the input sentence. We propose a novel

model that allows finer-grained control over what parts of the sentence are changed. Our approach leverages the sentence split categories described in §3.3 to identify the split-based edits and incorporates them into a customized loss function as distantly supervised labels. This section describes the base model and its variant that adapts a high paraphrasing BISECT corpus to a sentence splitting task with minimal rephrasing.

4.1 Base Model

Our base model is a **BERT-Initialized Transformer** (Rothe et al., 2020), a state-of-the-art model for Split and Rephrase. The encoder and decoder follow the BERT_{base} architecture, with the encoder initialized with the same checkpoint. The base model is trained using standard cross-entropy loss. During training, the split sentences in the reference are separated by a separator token [SEP].

Original Text	Split Text	WIKI	BiSECT
<i>Direct Insertion</i>		33%	40%
Gaal the son of Ebed came with his brothers, and went over to Shechem; and the men of Shechem put their trust in him. (fr→en)	Gaal the son of Ebed came with his brethren, and they passed over to Shechem. The people of Shechem trusted him.	<i>Colon/Semicolon</i> 15%	18%
When I play a MIDI file on my desktop, the sound quality is rich and clear, but when I play the same file on a laptop, it's not so great! (fr→en)	When I play MIDI files on my table extension the sound quality is excellent. If I play them on my portable sound is no longer very good.	<i>Conjunction with subject</i> 18%	22%
<i>Changes Near Split</i>		66%	49%
The virus is carried and passed to others through blood or sexual contact and can cause liver inflammation, fibrosis, cirrhosis and cancer. (de→en)	The virus is transmitted to other people through blood or sexual contact. It can cause liver inflammation, fibrosis, cirrhosis, and cancer.	<i>Conjunction without subject</i> 18%	13%
An additional advantage is that a shorter ramp can be used, thereby reducing weight and improving the rear view of the driver. (de→en)	Another advantage is that a shorter ramp can be used. This saves weight and improves the look of the rear of the vehicle.	<i>Gerund</i> 7%	10%
For the fur edge I choose the smudge tool with a dissolved brush and paint in the mask along the black edge to get a smooth transition. (de→en)	For the fur edge, I choose the tool with speckled brush tip and drag on the black edge in the mask. This creates a transition to the background .	<i>Preposition / Subordinate clause</i> 17%	9%
Over 3500 people visit the Centre every year where they are greeted by volunteers who show them around the study room and tell them about the collection. (fr→en)	Each year, more than 3,500 people visit the Center. They are greeted by volunteers who show them the study room and introduce them to the collection.	<i>Concluding Relative Clause</i> 24%	17%
<i>Changes Across Sentence</i>		1%	11% ↑
Because these cities, settlements and regions were constructed for not hundred years, but for centuries. (fr→en)	All these towns, these localities were not built in a hundred years. They were created over the centuries.	<i>Preceding Relative Clause</i> 1%	11%

Table 4: Categories in Split and Rephrase tasks with examples and frequency observed in the WIKISPLIT (Botha et al., 2018) and the English BiSECT (our work) corpora. Categories grouped under *Direct Insertion* require extremely minor changes in order to split the sentence; categories under *Changes Near Split* require some minor modifications around the source of the split; and categories under *Changes Across Sentence* require more major changes across the original sentence. Statistics are based on manual inspection of 100 examples from each corpus.

4.2 Adaptive Loss using Distant Supervision

The base model treats all the sentence splitting categories (Table 4) similarly even though the edits necessary to split the sentence vary across the categories. We utilize heuristics and linguistic rules to categorize each source-target sentence pair and extract required edits based on the category. Finally, we train the base model on these classification and edit labels to guide the model to perform appropriate edits for each category.

Classification and Edit Labels. Given the source $\mathbf{x} = (x_1, x_2, \dots, x_N)$ and target $\mathbf{y} = (y_1, y_2, \dots, y_N)$, we assign a sentence category label $l \in \{\text{"Direct Insertion"}, \text{"Changes Near Split"}, \text{"Changes Across Sentence"}\}$ to the training pair, and a binary label i to each position indicating whether the word is modified from the input. Here, $i = (i_1, i_2, \dots, i_N)$ represent the edit labels and $i_i = 1$ represents the necessary changes to split

the sentence that cannot be copied from \mathbf{x} . We ensure that \mathbf{x} and \mathbf{y} are of the same length using padding around the split. The split position for \mathbf{y} corresponds to the position of the [SEP] token. For \mathbf{x} , we extract the lexical differences between \mathbf{x} and \mathbf{y} using an edit distance algorithm⁵ and label the edit in \mathbf{x} close to the [SEP] token in \mathbf{y} as the split position. Finally, we pad the sequences before and after the split positions so that they are of equal length. We provide an example in Appendix D.

We extract l for each pair using the following rules: (1) If the first level of the parse tree of \mathbf{x} contains the pattern “S CC S”, \mathbf{x} contains a colon/semicolon, or the lexical differences between \mathbf{x} and \mathbf{y} contain only the split, then we label the pair as *Direct Insertion*. Once again, we extract lexical differences using an edit distance algorithm. (2) If the first level parse tree of \mathbf{x} contains the pat-

⁵<https://pypi.org/project/simplydiff/>

tern “*S NP VP*” or “*SBAR NP VP*”, then we label the pair as *Changes across sentence*. (3) If the first level of the parse tree contains “*VP CC VP*” or at least 5 words at the beginning and end of the sentence are copied from the source, then we categorize the pair as *Changes near split*. (4) We label the rest as *Changes across sentence*. In case of multiple potential splits, we choose the split whose lengths is closest to that of the reference.

After extracting l , we construct \hat{y}_i using the lexical overlap between x and y . For *Direct Insertion*, we set the i corresponding to the split position and its adjacent positions to 1 to capture the punctuation and capitalization. For *Changes near split*, we construct a variable length window around the split position to facilitate the addition of the new words and set the i in the window to 1. To construct this window, we scan the sequence on each side of the split position until the position where at least 3 consecutive positions are copied from x to y . Finally, we set \hat{y}_i to a one vector for *Changes Across Sentence*, as the changes cannot be localized. Our manual inspection of 100 training pairs from the BISECT training set showed that the rules correctly classified 83% of the pairs.

Distant Supervision. As l depends on the reference and cannot be used during inference, we introduce a multi-class classification task distantly supervised by l . We train our model in a multi-task learning setting to predict l and perform generation. The classifier predicts the probability that x belongs to a split category using the encoder representation of the $[CLS]$ token prepended to the input by the BERT encoder. The classifier contains a linear layer with a *softmax* activation function.

While l represents the sentence category, \hat{y}_i captures split-related edits. To ensure our model learns only split-based edits, we combine x and y in our decoder generation loss (L_{seq}) using \hat{y}_i as follows:

$$L_{seq} = \frac{1}{m} \sum_{i=1}^m (1 - \hat{y}_i) P(x_i | \hat{y}_{<i}) + \hat{y}_i P(y_i | \hat{y}_{<i})$$

$$\hat{y}_i = (1 - \hat{y}_i) x_i + \hat{y}_i y_i$$

$$\hat{y}_i = \begin{cases} 0; & \text{if } x_i \text{ is copied} \\ 1; & \text{otherwise} \end{cases}$$

where m is the number of training examples and $\hat{y}_{<i}$ represents the mixture of x and y histories. In other words, our model only learns the edits where $\hat{y}_i = 1$ and copies from the source sentence for the rest of the positions. Finally, we jointly train

the classifier and the Transformer using the cross entropy loss and our custom split-focused loss. We provide model and training details in Appendix A.

5 Experiments and Results

In this section, we compare different split and rephrase models trained on our new BISECT corpus. We also conduct a carefully designed human evaluation as automatic metrics are not totally reliable. Our model trained on BISECT establishes a new start-of-the-art for the task.

5.1 Data and Baselines

We train the models on BISECT and WIKISPLIT corpora. For evaluation, we select the BISECT and HSPLIT-WIKI (Sulem et al., 2018) test sets to represent splitting with a high degree and minimal of rephrasing respectively. HSPLIT-WIKI is a human annotated dataset with 359 complex sentences and 4 references for each complex sentence. Following previous work (Botha et al., 2018; Zhang et al., 2020a), we do not use WIKISPLIT for evaluation, because this corpus was constructed explicitly to be used only as training data, as it contains inherent noise and biases. While BISECT contains 928,440/9,079 train and dev pairs, WIKISPLIT contains 989,944/5,000 train and dev pairs. Note that we constructed BISECT test set by manually selecting 583 high-quality sentence splits from 1000 random source-target pairs from EMEA and JRC-ACQUIS corpora.

We compare our approach with **Copy512** (Aharoni and Goldberg, 2018), a state-of-the-art model consisting of an attention-based LSTM encoder-decoder with a copy mechanism (See et al., 2017). We use our base model trained on WIKISPLIT (Rothe et al., 2020) as another state-of-the-art baseline.

5.2 Automatic Evaluation

Existing automatic metrics, such as BLEU (Papineni et al., 2002) and SAMSA (Sulem et al., 2018), are not optimal for the Split and Rephrase task as they rely on lexical overlap between the output and the target (or source) and underestimate the splitting capability of the models that rephrase often. We focus on BERTScore (Zhang et al., 2020b) and SARI (Xu et al., 2016). BERTScore (Zhang et al., 2020b) captures meaning preservation and fluency well (Scialom et al., 2021). SARI can provide three separate F1/precision scores that explicitly mea-

Models w/ Training Data	SARI	add	keep	del	BScore	FK	BLEU	SLen	OLen	sBLEU	%new
BiSECT test set											
Source	20.1	0.0	60.3	0.0	84.6	19.2	43.5	35.2	39.9	100.0	0.0
DisSim	40.0	2.4	55.2	62.4	76.9	11.2	30.0	12.4	40.6	61.4	18.7
Copy512 w/ WIKI	46.7	4.0	61.6	74.6	84.1	13.0	43.0	19.4	39.7	88.5	4.5
Copy512 w/ BiSECT	52.7	10.6	64.8	82.8	85.3	12.6	46.3	18.5	39.2	81.5	6.7
Transformer w/ WIKI	49.3	6.9	62.8	78.2	84.5	12.4	43.1	19.3	41.0	81.8	9.3
Transformer w/ BiSECT	55.5	18.3	66.9	81.4	85.6	12.1	45.8	19.0	40.7	63.9	16.6
Transformer _{control} w/ BiSECT	49.0	7.9	62.6	76.1	84.8	12.6	42.9	19.8	40.9	79.7	10.5
Transformer _{control} w/ BiSECT+WIKI	47.7	6.0	62.1	75.0	84.3	12.9	43.5	19.1	39.7	85.5	6.0
Reference	94.3	88.8	97.9	96.1	100.0	12.5	100.0	19.2	41.5	40.4	32.0
HSPLIT-WIKI test set											
Source	30.5	0.0	91.4	0.0	97.1	12.6	83.0	22.4	22.6	100.0	0.0
DisSim	38.0	5.0	79.3	29.6	87.7	8.9	52.5	10.5	23.6	62.8	17.1
Copy512 w/ WIKI [†]	47.2	13.0	87.9	40.8	93.3	8.4	68.2	12.3	24.7	71.2	17.0
Copy512 w/ BiSECT	47.4	13.6	87.4	40.7	92.3	8.3	69.0	12.0	23.6	72.2	14.3
Transformer w/ WIKI [†]	49.5	14.9	88.4	45.2	95.3	7.8	69.2	12.0	24.8	73.1	15.8
Transformer w/ BiSECT	45.7	17.7	80.2	39.1	92.0	7.8	57.8	12.7	26.2	57.0	26.2
Transformer _{control} w/ BiSECT	47.2	13.3	87.2	41.1	94.1	7.9	67.2	12.3	24.9	70.9	17.6
Transformer _{control} w/ BiSECT+WIKI	52.0	15.7	90.4	50.0	95.4	8.3	74.0	11.9	23.9	78.2	11.9
Reference	60.1	33.0	94.1	53.2	100.0	8.4	100.0	12.6	24.3	81.8	10.6

Table 5: Automatic and human evaluation results on BiSECT and HSPLIT-WIKI test sets. We report **SARI** and its three edit scores, namely precision for delete (**del**) and F1 scores for **add** and **keep** operations. We also report BERTScore (**BScore**), FKGL (**FK**), corpus-level BLEU (**BLEU**), average number of words in a sentence (**SLen**), average number of words in the output (**OLen**), self-BLEU (**sBLEU**), and average percentage of new words added to the output (**%new**). **Bold** typeface denotes the best performances (i.e., closest to the reference). [†] These models have a natural advantage on the i.i.d. sampled Wiki-based HSPLIT test set, as they are trained on WIKISPLIT data. In contrast, the train and test data in BiSECT are not i.i.d. sampled and from different sources (Table 2).

Model	BiSECT	HSPLIT-WIKI
Random	13.2	11.9
Transformer w/ WIKI	88.2	88.1
Transformer w/ BiSECT	93.8	92.0
Transformer _{control} w/ B	94.8	84.5
Transformer _{control} w/ B+W	89.2	88.5
Reference	95.0	96.8

Table 6: Human evaluation of the overall sentence splitting quality (rating on 0-100 scale) on 100 examples from the BiSECT and HSPLIT-WIKI test sets. **B** and **W** represent BiSECT and WIKISPLIT respectively.

sure the correctness of inserted, kept and deleted n-grams when compared to both the source and the target. We use an extended version of SARI that considers lexical paraphrases of the reference. An n-gram from the output is considered correct if the given n-gram or its paraphrase from PPDB (Pavlick et al., 2015) occurs in the reference, using the PPDB-L version. Without this change, the original SARI also tends to underestimate rephrasing.

Table 5 shows that our models trained on BiSECT outperform their equivalents trained on WIKISPLIT in terms of SARI and BERTScore. Note that the models trained on WIKISPLIT have an advantage over HSPLIT-WIKI test set because they belong to the same domain. Models trained

on BiSECT do not have a similar advantage on BiSECT test set because it belongs to a different domain than the training data. When compared to the base model (Transformer w/ BiSECT), our model (Transformer_{control} w/ BiSECT) shows higher self-BLEU and lower percentage of new words, indicating that it performs less rephrasing by focusing on split-based edits.

5.3 Human Evaluation

We asked three annotators to rate the overall quality of the sentence splits generated by different models on a 0-100 point scale. 0 represents an erroneous split and 100 represents a perfect meaning-preserving split. Unlike the previous work that measures meaning preservation and fluency separately, we collected only one rating because it was difficult to distinguish between the grammatical and the meaning-changing errors. We modeled our evaluation after the WMT evaluation (Bojar et al., 2019) that also uses a similar setting. We evaluated on 100 random sentences from the BiSECT and HSPLIT-WIKI test sets. The annotators were university students trained using an instructional video and a qualification phase. To capture the annotation quality, we included a control output generated by

Model w/ Data	System Outputs
Source	Having determined, after consulting the Advisory Committee that sufficient evidence existed for the initiation of a partial interim review, the Commission published a notice in the Official Journal of the European Communities and commenced an investigation.
Transformer w/ WIKI	Having determined, after consulting the Advisory Committee, that sufficient evidence existed for the initiation of a partial interim review. The Commission published a notice in the Official Journal of the European Communities and commenced an investigation.
Transformer w/ BI-SECT	After consulting the Advisory Committee, the Commission determined that there was sufficient evidence for the initiation of a partial interim review. The Commission issued a notice in the Official Journal of the European Communities and began an investigation.
Transformer _{control} w/ BI-SECT	After consulting the Advisory Committee, there was sufficient evidence for the initiation of a partial interim review. The Commission published a notice in the Official Journal of the European Communities and initiated an investigation.

Table 7: Examples of system outputs from the BI-SECT test set. Here, the source sentence belongs to the category “Changes Across Sentence”. Blue marks the location of the required edits in the source sentence. Green indicates good edits and red indicates errors.




Figure 2: Human ratings on 100 generated sentence splits from the BI-SECT test set broken down by sentence split categories as described in Table 4.

randomly selecting a system output and replacing 4 to 8 words with random words. Our annotators gave low ratings (<20) to the control outputs, indicating that the ratings are reliable. We provide the annotation interface design in Appendix F.

Table 6 shows that results on the entire BI-SECT and WIKI-HSPLIT test sets. Figure 2 shows the results on different split categories of the BI-SECT test set. The sentences splits generated by models trained on BI-SECT are of better quality than the ones trained on WIKISPLIT. Our model with adaptive loss (Transformer_{control} w/ BI-SECT) performs better than the base model (Transformer w/ BI-SECT) in four of the seven split categories. The difference in quality is much more evident for the *Preceding Relative Clause* category, as this requires changes across sentences. We provide an example in Table 7, as well as several more in Appendix E.

6 Conclusion

In this work, we introduce BI-SECT, a new corpus for the Split and Rephrase task in several languages. We create this by making use of bilingual parallel corpora, and translating instances of aligned split sentences. We show that the sentence splitting

models trained on our new corpus generate fewer errors than their counterparts trained on the existing datasets. To further improve meaning preservation and diversity, we propose a novel approach that identifies split-related edits in a training pair using linguistic rules and trains the model solely on split-based edits. Our proposed approach trained on BI-SECT outperforms existing systems in terms of both automatic and human evaluations. We plan to investigate and create better automatic evaluation metrics for future work.

Acknowledgments

We thank four anonymous reviewers for their helpful comments. We also thank Andrew Duffy, Manish Jois, Riley Kolb, Sounak Dey, Alyssa Hwang, Bryan Li, Veronica Qing Lyu and UPenn NETS 212 students for assisting through the human evaluation iterations. This research is supported in part by the NSF awards IIS-2055699, ODNI and IARPA via the BETTER program contract 19051600004, and the DARPA KAIROS Program (contract FA8750-19-2-1004). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of NSF, ODNI, IARPA, DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

Roe 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, pages 719–724.
- Colin Bannard and Chris Callison-Burch. 2005. Paraphrasing with bilingual parallel corpora. In Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05), pages 597–604.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Mark Fishel, Yvette Graham, Barry Hadrow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, André Martins, Christof Monz, Matteo Negri, Aurélie Névéal, Mariana Neves, Matt Post, Marco Turchi, and Karin Verspoor, editors. 2019. Proceedings of the Fourth Conference on Machine Translation.
- 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.
- Michael Collins, Philipp Koehn, and Ivona Kučerová. 2005. Clause restructuring for statistical machine translation. In Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05), pages 531–540.
- William A. Gale and Kenneth W. Church. 1993. A program for aligning sentences in bilingual corpora. Computational Linguistics, 19(1):75–102.
- Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2013. PPDB: The paraphrase database. In Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 758–764.
- 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.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. Creating Training Corpora for NLG Micro-Planning. In 55th annual meeting of the Association for Computational Linguistics (ACL).
- J Edward Hu, Rachel Rudinger, Matt Post, and Benjamin Van Durme. 2019a. Parabank: Monolingual bitext generation and sentential paraphrasing via lexically-constrained neural machine translation. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, pages 6521–6528.
- J. Edward Hu, Abhinav Singh, Nils Holzenberger, Matt Post, and Benjamin Van Durme. 2019b. Large-scale, diverse, paraphrastic bitexts via sampling and clustering. In Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL), pages 44–54.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In 3rd International Conference on Learning Representations, ICLR 2015.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In Proceedings of the First Workshop on Neural Machine Translation, pages 28–39.
- Mounica Maddela, Fernando Alva-Manchego, and Wei Xu. 2021. Controllable text simplification with explicit paraphrasing. In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 3536–3553, Online. Association for Computational Linguistics.
- Jonathan Mallinson, Rico Sennrich, and Mirella Lapata. 2017. Paraphrasing revisited with neural machine translation. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, pages 881–893.
- Jonathan Mallinson, Rico Sennrich, and Mirella Lapata. 2018. Sentence compression for arbitrary languages via multilingual pivoting. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 2453–2464.
- Ryan McDonald and Joakim Nivre. 2011. Analyzing and integrating dependency parsers. Computational Linguistics, 37(1):197–230.
- 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, pages 435–445.
- 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.
- 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.
- Kostiantyn Omelianchuk, Vitaliy Atrasevych, Artem Chernodub, and Oleksandr Skurzhanskyi. 2020. GECToR – grammatical error correction: Tag, not rewrite. In Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications, pages 163–170, Seattle, WA, USA. Online. 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.
- Ellie Pavlick, Pushpendre Rastogi, Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2015. [PPDB 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification](#). In [Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing](#), pages 425–430.
- Jean Pouget-Abadie, Dzmitry Bahdanau, Bart van Merriënboer, Kyunghyun Cho, and Yoshua Bengio. 2014. [Overcoming the curse of sentence length for neural machine translation using automatic segmentation](#). In [Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation](#), pages 78–85.
- Sascha Rothe, Shashi Narayan, and Aliaksei Severyn. 2020. [Leveraging pre-trained checkpoints for sequence generation tasks](#). [Transactions of the Association for Computational Linguistics](#), 8:264–280.
- Thomas Scialom, Louis Martin, Jacopo Staiano, Éric Villemonte de la Clergerie, and Benoît Sagot. 2021. [Rethinking automatic evaluation in sentence simplification](#).
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In [Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics](#), pages 1073–1083.
- Elior Sulem, Omri Abend, and Ari Rappoport. 2018. [BLEU is not suitable for the evaluation of text simplification](#). In [Proceedings of EMNLP 2018 Conference on Empirical Methods in Natural Language Processing](#), pages 738–744.
- Jörg Tiedemann and Lars Nygaard. 2004. [The OPUS corpus - parallel and free: <http://logos.uio.no/opus>](#). In [Proceedings of the Fourth International Conference on Language Resources and Evaluation \(LREC'04\)](#).
- Mark Twain. 1880. [The awful German language](#). BVK.
- John Wieting and Kevin Gimpel. 2018. [ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In [Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics](#), pages 451–462.
- Wei Xu, Chris Callison-Burch, and Courtney Napoles. 2015. [Problems in current text simplification research: New data can help](#). [Transactions of the Association for Computational Linguistics](#), 3(1):283–297.
- 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(1):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 \(EMNLP\)](#).
- 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.

A Implementation and Training details

We implemented the BERT-initialized Transformer using the Fairseq⁶ toolkit. Here, the encoder and decoder follow BERT_{base}⁷ architecture. The encoder is also initialized with BERT_{base} checkpoint and the decoder is randomly initialized. The sentence classifier is a feedforward network containing an input layer, one hidden layer with 1000 nodes, and an output layer with 3 nodes and *softmax* activation. We used Adam optimizer (Kingma and Ba, 2015) with a learning rate of 0.0001, linear learning rate warmup of 40k steps, and 100k training steps. We used a batch size of 64. We used BERT WordPiece tokenizer. During inference, we use beam-search of width 10 and ensure that the beam-search does not repeat trigrams. We used the hyperparameters of the BERT-initialized Transformer described in Rothe et al. (2020). The model takes 10 hours to train on 1 NVIDIA GeForce GPU.

⁶<https://github.com/pytorch/fairseq>

⁷<https://github.com/google-research/bert>

B BiSECT Language Composition

Dataset	French	German	Spanish	Arabic	Dutch	Italian	Portuguese	Russian
CCALIGNED	204K	–	–	–	–	–	–	–
EUROPARL	57K	–	–	–	–	–	–	–
GIGAWORD	132K	–	–	–	–	–	–	–
PARACRAWL	125K	144K	66K	–	31K	26K	14K	–
UN	5.7K	–	8K	36K	–	–	–	6.7K
EMEA	1K	–	–	–	–	–	–	–
JRC-ACQUIS	1K	3.7K	–	–	–	–	–	–
TOTAL	672K	151K	75K	36K	33K	27K	15K	6.7K

Table 8: Composition of the BiSECT (English) corpus by pivoted language over bilingual parallel corpora.

C Examples from Different Corpora

Corpus	Examples
WIKI-AUTO	<p>Source: Following the establishment of the Pembrokeshire Coast National Park in 1952, Welsh naturalist and author Ronald Lockley surveyed a route around the coast.</p> <p>Reference: The Pembrokeshire Coast National Park was founded in 1952. After it was founded, Ronald Lockley did a survey for a path on the coastline.</p>
NEWSELA-AUTO	<p>Source: About 160,000 Girl Scouts participated in the program over the past year and were credited with selling nearly 2.5 million boxes of cookies beyond those sold through traditional in-person methods.</p> <p>Reference: About 160,000 Girl Scouts used Digital Cookie last year. They sold almost 2.5 million boxes of cookies online.</p>
HSPLIT	<p>Source: West Berlin had its own postal administration, separate from West Germany’s, which issued its own postage stamps until 1990.</p> <p>Reference: West Berlin had its own postal administration. It was separate from West Germany’s. West Berlin issued its own postage stamps until 1990.</p>
CONTRACT	<p>Source: Except for Supplier’s obligations and liability resulting from Section 10.0, Supplier Liability for Third Party Claims, Supplier’s liability for any and all claims will be limited to the amount of \$1,000,000 USD per occurrence, with an aggregated limit of \$4,500,000 USD during the term of this Agreement.</p> <p>Reference: The following applies, not including the Supplier’s obligations and liability resulting from Section 10.0, Supplier Liability for Third Party Claims. Supplier’s liability for any and all claims will be limited to the amount of \$1,000,000 USD per occurrence. Additionally, there is an aggregated limit of \$4,500,000 USD during the term of this Agreement.</p>
WIKI-BM	<p>Source: Together with James, she compiled crosswords for several newspapers and magazines, including People, and it was in 1978 that they launched their own publishing company.</p> <p>Reference: Together with James, she compiled crosswords. It was for several newspapers and magazines, including People. They launched their own publishing company. It was in 1978.</p>
WEBSPLIT v1.0	<p>Elliot See (born on July 23, 1927 in Dallas and died on February 28, 1966 in St Louis) was an American who graduated from the University of Texas at Austin.</p> <p>Elliot See attended the University of Texas at Austin. Elliot See, deceased, was born in Dallas. Elliot See died on February 28, 1966, in St Louis. Elliot See was born on July 23, 1927. Elliot See is a United States national.</p>
WIKISPLIT	<p>In 2006, he and the Cavaliers negotiated a three-year, \$ 60 million contract extension instead of the four year maximum as it allotted him the option of seeking a new contract worth more money as an unrestricted free agent following the 2010 season.</p> <p>In 2006, he and the Cavaliers negotiated a three-year, \$ 60 million contract extension. This was instead of the four year maximum length as it allotted James the option of seeking a new contract worth more money as an unrestricted free agent following the 2010 season.</p>
BiSECT	<p>Respondents felt that headsets compatible with hearing aids would greatly assist them in understanding what is being said, and added that headsets in business class or first class on some aircraft are already compatible with hearing aids.</p> <p>Participants indicated that the installation of headsets which are compatible with hearing aids would improve their ability to understand what was being said. It was mentioned that headsets in the business or first class portions of some aircraft are already hearing aid compatible.</p>

Table 9: Random examples of sentence pairs from the existing corpora. Blue indicates the position of sentence splits in the source sentence. Green indicates good edits, and red indicates hallucinations in the reference.

D Our Model

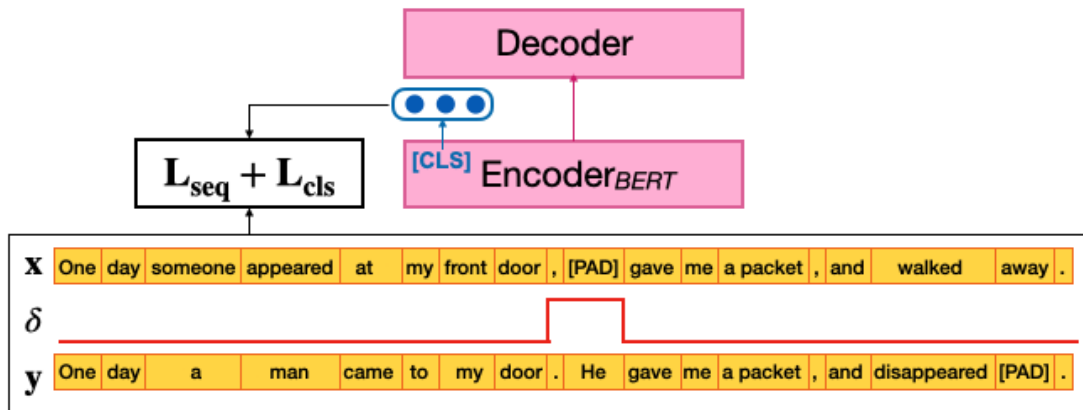


Figure 3: Overview of our proposed approach, where the model is trained on only the split-related edits in y determined by edit and sentence category l labels. We also illustrate the padding and the construction of δ for a training pair belonging to *Changes Near Split* category.

E Examples of System Outputs

Model w/ Data	System Outputs
<i>Preposition / Subordinate Clause</i>	
Source	To that end, the annual contribution from each Member State shall be established on the basis of the gross national product (GNP) according to the scale used for determining the GNP element in own resources for financing the general budget of the European Union.
Transformer w/ WIKI	To that end, the annual contribution from each Member State shall be established on the basis of the gross national product (GNP). According to the scale used for determining the GNP element in own resources for financing the general budget of the European Union .
Transformer w/ BISECT	To this end, the annual contribution of each Member State shall be established on the basis of the gross national product (GNP). the scale used to determine the GNP element in own resources for the financing of the general budget of the European Union .
Transformer _{control} w/ BISECT	To that end, the annual contribution from each Member State shall be established on the basis of the gross national product. This depends on the scale used to determine the GNP element in own resources for financing the general budget of the European Union .
<i>Gerund</i>	
Source	Member States shall take appropriate and practicable measures to reduce the risk of transmission of avian influenza from birds living in the wild to susceptible birds kept in zoos, taking into account the criteria and risk factors set out in Annex I.
Transformer w/ WIKI	Member States shall take appropriate and practicable measures to reduce the risk of transmission of avian influenza from birds living in the wild to susceptible birds kept in zoos. Taking into account the criteria and risk factors set out in Annex I.
Transformer w/ BISECT	Member States shall take appropriate and practical measures to reduce the risk of transmission of avian influenza from birds living in the wild to susceptible birds kept in zoos. These measures will take into account the criteria and risk factors set out in Annex I.
Transformer _{control} w/ BISECT	Member States shall take appropriate and practicable measures to reduce the risk of transmission of avian influenza from birds living in the wild to susceptible birds kept in zoos. They must take into account the criteria and risk factors set out in Annex I.
<i>Concluding Relative Clause</i>	
Source	There were two main studies involving a total of 529 patients with SLE, in which Riquent was compared with a placebo (a dummy treatment).
Transformer w/ WIKI	There were two main studies involving a total of 529 patients with SLE. In which Riquent was compared with a placebo (a dummy treatment).
Transformer w/ BISECT	There were two main studies involving a total of 529 patients with SLE. Riquent was compared with a placebo (a dummy treatment).
Transformer _{control} w/ BISECT	There were two main studies involving a total of 529 patients with SLE. Riquent was compared with a placebo (a dummy treatment).
<i>Conjunction Without Subject</i>	
Source	Undertakings required to maintain separate accounts are undertakings that enjoy special or exclusive rights granted by a Member State as referred to in Article 86 (1) of the Treaty or are entrusted with the operation of a service of general economic interest as referred to in Article 86 (2) of the Treaty and receive State aid in any form whatsoever in relation to such service and that carry on other activities .
Transformer w/ WIKI	Undertakings required to maintain separate accounts are undertakings that enjoy special or exclusive rights granted by a Member State as referred to in Article 86 (1) of the Treaty. They may also be entrusted with the operation of a service of general economic interest as referred by the Treaty and receive State aid in any form whatsoever in relation to such service and that carry on other activities .
Transformer w/ BISECT	Undertakings required to maintain separate accounts are undertakings that enjoy special or exclusive rights granted by a Member State under Article 86 (1) of the Treaty or are entrusted with the operation of a service of general economic interest referred to in Article 86 Para . 2) and receive State aid in any form in relation to this service and carry out other activities .
Transformer _{control} w/ BISECT	Undertakings required to maintain separate accounts are undertakings that enjoy special or exclusive rights granted by a Member State as referred to in Article 86 (1) of the Treaty or are entrusted with the operation of a service of general economic interest. They must also receive state aid in any form whatsoever in relation to such service and that carry on other activities .
<i>Preceding Relative Clause</i>	
Source	Because 1'-hydroxymidazolam is an active metabolite, the sedative effect of midazolam may be increased.
Transformer w/ WIKI	Because 1'- hydroxymidazolam is an active metabolite. The sedative effect of midazolam may be increased.
Transformer w/ BISECT	1'-hydroxymidazolam is an active metabolite. The sedative effect of midazolam can be increased.
Transformer _{control} w/ BISECT	1'-hydroxymidazolam is an active metabolite. The sedative effect of midazolam may therefore be increased.

Table 10: Examples of system outputs from the BISECT test set. Here, the source sentence belongs to the category “Changes Near Split”. Blue marks the location of the required edits in the source sentence. Green indicates good edits and red indicates errors.

F Human Evaluation

Quality	Rate Split Quality				English
---------	--------------------	--	--	--	---------

Read the following pairs of texts (source and candidate) and provide ratings between 0 to 100 based on the degree of similarity in **meaning** and preservation of **grammar** for each candidate text. Drag or click on the appropriate portion of the slider to provide a rating for each candidate text. Please refer to the scoring example before starting the first task.

Note: There will be a candidate text in each HIT which contains several random words appearing out of context unrelated to the text. Please make sure to score this candidate text in each HIT with **10-20** points. Apart from this, please refer to the table below for scoring the candidate text. You must provide a rating for all examples to proceed.

Category	Mild (1 sent)	Significant (1 sent)	Mild (2 sent)	Significant (2 sent)
Meaning	deduct 5-10 pts e.g., wrong pronoun	deduct 30 pts e.g., relevant words but diff meaning	deduct 10-20 pts	deduct 60-70 pts
Grammaticality	deduct 5-10 pts e.g., missing conjunction	deduct 25 pts e.g., incomplete sentence	deduct 10-20 pts	deduct 50-60 pts

Hide/Show Examples

The vendor will expertly cut the top from the nut, insert a straw and, voila, you will have a hygienic, refreshing, healthful drink.

The [^]seller will [^]cut the top from the nut, insert a straw [^]and voila. You will have a hygienic, [^]refreshing and healthy drink.

Very Low Quality (0) Perfect Quality (100)

The [^]seller will [^]skillfully cut the top [^]of the [^]fruit, insert a [^]straw, and voila! you will have [^]an uncontaminated, refreshing and healthy drink.

Very Low Quality (0) Perfect Quality (100)

The [^]seller will expertly cut the top [^]of the [^]nut and insert a [^]straw. In addition, you will have a hygienic, [^]refreshing and healthy drink...

Very Low Quality (0) Perfect Quality (100)

The vendor will expertly cut the top from the nut, insert a straw and, [^]voila. You will have a hygienic, refreshing, healthful drink.

Very Low Quality (0) Perfect Quality (100)

The [^]seller will [^]skillfully motor the top [^]of the [^]fall insert a [^]waning whiter voila! you will have [^]an uncontaminated, refreshing and healthy drink.

Very Low Quality (0) Perfect Quality (100)

Reset Submit

Figure 4: Annotation interface and guidelines for human evaluation. Each system output is followed by a slider ranging between 0 to 100 with labels “*Very Low Quality*” on the left and “*Perfect Quality*” on the right. Highlighted words indicate newly added words when compared to the source sentence. Hovering the mouse over the red ticks displays words removed from the source sentence. Every HIT contains a control text, where 4 to 8 words are replaced with random words. Workers are expected to give low scores to the control text. Furthermore, the system outputs are shuffled for every HIT to eliminate position bias.

G Multilingual BiSECT

G.1 French

Dataset	Pivot Lang.	Domain	1-2 & 2-1 Alignments		Sent. Length	
			total (count/%)	after filtering	long	split
CCALIGNED	en	web crawl	164,628 (12.85%)	56,799	37.16	41.05
EUROPARL	en	European Parliament	153,220 (11.96%)	57,581	46.30	47.74
GIGAWORD	en	newswire	624,372 (48.73%)	235,133	43.73	44.58
PARACRAWL	en	web crawl	308,047 (24.04%)	127,655	39.07	39.03
UN	en	United Nations	23,706 (1.85%)	13,869	47.45	49.93
EMEA	en	European Medicines Agency	5,719 (0.45%)	2,400	40.03	45.12
JRC-ACQUIS	en	European Union	1,690 (0.13%)	1,036	49.11	52.94

Table 11: Statistics of datasets in the OPUS collection that we used to create French version of the BiSECT corpus.

G.2 Spanish

Dataset	Pivot Lang.	Domain	1-2 & 2-1 Alignments		Sent. Length	
			total (count/%)	after filtering	long	split
CCALIGNED	en	web crawl	466,240 (56.16%)	110,958	40.45	46.11
PARACRAWL	en	web crawl	297,879 (35.88%)	162,048	35.36	33.75
UN	en	United Nations	17,948 (2.16%)	9,938	48.02	51.76
EUROPARL	en	European Parliament	48,165 (5.80%)	6,719	46.68	47.90

Table 12: Statistics of datasets in the OPUS collection that we used to create Spanish version of the BiSECT corpus.

G.3 German

Dataset	Pivot Lang.	Domain	1-2 & 2-1 Alignments		Sent. Length	
			total (count/%)	after filtering	long	split
CCALIGNED	en	web crawl	510,817 (52.57%)	52,253	30.87	36.65
EUROPARL	en	European Parliament	100,784 (10.37%)	16,359	42.08	44.24
PARACRAWL	en	web crawl	353,136 (36.34%)	116,026	37.73	38.99
JRC-ACQUIS	en	European Union	6,950 (0.72%)	1,599	54.79	55.19

Table 13: Statistics of datasets in the OPUS collection that we used to create German version of the BiSECT corpus.