

# Does mBERT Understand Romansh? Evaluating Word Embeddings Using Word Alignment.

Eyal Liron Dolev

Zurich Center of Linguistics – LiZZ

University of Zurich

eyalliron.dolev@uzh.ch

## Abstract

We test similarity-based word alignment models (SimAlign and awesome-align) in combination with word embeddings from mBERT and XLM-R on parallel sentences in German and Romansh. Since Romansh is an unseen language, we are dealing with a zero-shot setting. Using embeddings from mBERT, both models reach an alignment error rate of 0.22, which outperforms fast\_align, a statistical model, and is on par with similarity-based word alignment for seen languages. We interpret these results as evidence that mBERT contains information that can be meaningful and applicable to Romansh.

To evaluate performance, we also present a new trilingual corpus, which we call the DERMIT (DE-RM-IT) corpus, containing press releases made by the Canton of Grisons in German, Romansh and Italian in the past 25 years. The corpus contains 4,547 parallel documents and approximately 100,000 sentence pairs in each language combination. We additionally present a gold standard for German-Romansh word alignment. The data is available at <https://github.com/eyldlv/DERMIT-Corpus>.

## 1 Introduction

Romansh is a Romance language spoken by an ever-diminishing number of speakers (approximately 40,000 as of 2020<sup>1</sup>) mostly in the Canton of Grisons (henceforth Graubünden). The name *Romansh* is an umbrella term for five dialects (Surselvan, Sutselvan, Surmiran, Puter, Vallader), each having normative grammars and orthography (Haiman and Benincà, 1992, p. 1). In 1982, Rumantsch Grischun, a standard written variety of

Romansh was launched, its goal being the harmonization of the five dialects and the creation of a language variant that would be understood by all speakers of Romansh and could be used i.a. by officials, government agencies and in formal writings (Haiman and Benincà, 1992, p. 5).<sup>2</sup>

Currently, there are no state-of-the-art NLP tools, such as part-of-speech (PoS) tagging or named entity recognition (NER), for Romansh. However, multilingual language models (MLMs) were shown to perform well in various tasks in zero-shot settings and on unseen languages, i.e., for languages not included in the training data. Pires et al. (2019) showed that fine-tuning mBERT for PoS-tagging on one language generalizes well (over 80% accuracy) for other languages. Similarly, the LASER model, which was pre-trained on 93 languages, obtained strong results for sentence embeddings in 112 languages (Artetxe and Schwenk, 2019). Since Romansh data, raw as well as annotated, is scarce, leveraging MLMs for annotation tasks and other tasks requiring word embeddings, could provide significant advantages when developing language technology and NLP applications for Romansh, instead of training new models from scratch.

We hypothesize that due to the proven success of MLMs in zero-shot settings and on unseen languages, and due to the similarity of Romansh to other Romance languages that were included in pre-training, MLMs should contain information that is meaningful and applicable to processing Romansh. To support this hypothesis, two conditions should be fulfilled:

Word alignments computed with similarity-based systems (such as SimAlign or awesome-align) using word embeddings taken from MLMs should:

1. be on par with statistical models;

<sup>1</sup><https://www.bfs.admin.ch/bfs/de/home/statistiken/bevoelkerung/sprachen-religionen/sprachen.assetdetail.21344032.html>

<sup>2</sup>In our paper, whenever we use the name *Romansh* we refer to the standard variety *Rumantsch Grischun*

2. be on par with alignments computed using multilingual embeddings for seen language pairs.

If both conditions are fulfilled, we would view this as an indication that these models contain information that can be meaningfully applied to processing Romansh.

To test our hypothesis, we collected a new trilingual corpus, which is a continuation of the ALLEGRA corpus by Scherrer and Cartoni (2012) (Section 3), extracted parallel sentences (Section 3.2) and created a word alignment gold standard by annotating 600 sentences (Section 4). The experiments and the results are described in Section 5.

## 2 Previous Work

Romansh has enjoyed a considerable amount of attention from the NLP community despite its low number of speakers. This is arguably due to several favorable circumstances: It is spoken in a highly modernized land, it enjoys official legal status as a national language and it is promoted and protected by law.<sup>3</sup>

Scherrer and Cartoni (2012) compiled a trilingual corpus using the press releases published by the canton of Graubünden, Baumgartner et al. (2014) created a system for automatic morphological analysis, Weibel (2014) compiled two bilingual (German-Romansh), sentence- and word-aligned corpora, containing the cantonal press releases and legal texts. These corpora are available on the website “multilingwis”<sup>4</sup> (Graën et al., 2017). Romansh was also part of the multilingual corpus used by Volk and Clematide (2014) for evaluating code-switching detection and it was used by Müller et al. (2020) for evaluating the performance of out-of-domain machine translation. Last year, in 2022, TextShuttle, a Zurich-based company specializing in machine translation, developed and released a machine translation system for Romansh (translating to or from German, French, Italian and English).

## 3 Corpus

Our corpus is based on public press releases published by the Canton of Graubünden.

<sup>3</sup>Cf. *Verfassung des Kantons Graubündens* and *Sprachengesetz des Kantons Graubündens*.

<sup>4</sup><https://pub.cl.uzh.ch/projects/sparcling/multilingwis2.demo/>

The Canton of Graubünden is the only Swiss canton with three official languages: German, spoken mostly in the northern parts of the canton, Italian, spoken mostly in the south, e.g., in the valleys Poschiavo and Mesolcina, and Romansh, spoken sporadically in several valleys, e.g., Engadine or Domleschg.

The canton has been publicly publishing press releases in these three official languages (for Romansh, the standard variety “Rumantsch Grischun” is used) at least since 1997, and they are publicly available on the canton’s website. The press releases are short texts, usually consisting of about 3-5 paragraphs, which are intended to inform the public about events and occurrences in the canton and in the cantonal authorities, about decisions taken by the cantonal government, etc. Some typical topics are politics, economics and health.

### 3.1 Scraping and Document Alignment

To collect said press releases and create a parallel corpus, we collected all the press releases published from 2010 until January 2023 on the canton’s website. We then aligned the documents using URL matching. The press releases published between 1997 and 2009 cannot be aligned automatically using URL matching. However, since Scherrer and Cartoni (2012) already manually aligned these documents for their ALLEGRA corpus, we included them, as is, in our own corpus. We thus see our work as a continuation of the ALLEGRA corpus.

Table 3 in Appendix A gives information about the total number of documents in each language and the number of documents that exist in at least two languages.

Both the aligned and the unaligned documents are presented as JSON files (one file per year) and as an SQLite database, which both allow for simple and fast queries.

### 3.2 Sentence Alignment

For sentence tokenization/segmentation, we used NLTK’s (Bird et al., 2009) Punkt tokenizer, which is pre-trained for German and Italian. For Romansh, we used the Italian tokenizer, which we extended with a list of Romansh abbreviations to improve sentence segmentation.

We then sentence-aligned the press releases using hunalign (Varga et al., 2005), a hybrid sentence-alignment system combining a dictionary-based and a length-based method. Although not a requirement, we supplied hunalign with a German-

Comb.	Sentences	Tokens	Types
DE-RM	106,091	1.9M, 2.4M	97K, 51K
DE-IT	103,441	1.8M, 2.2M	96K, 58K
IT-RM	102,757	2.3M, 2.2M	50K, 58K

Table 1: Unique sentence pairs extracted from the corpus for the three language combinations: German-Romansh, German-Italian and Italian-Romansh.

Romansh dictionary<sup>5</sup>. We have no gold standard for sentence alignment, but out of 611 sentences extracted for the task of annotating the gold standard, only 11 sentences had to be removed due to misalignment, which corresponds to an alignment error rate of 1.8%.

After removing duplicates and obvious misalignments using simple heuristics, we were able to extract 106,091 German-Romansh sentence pairs, 103,441 German-Italian sentence pairs and 102,757 Romansh-Italian sentence pairs, see also Table 1.

### 3.3 Word Alignment

Word alignments are objects mapping each word in a string of the source language  $e$  to the words in a string of the target language  $f$  (Koehn, 2009, p. 84). They are traditionally computed using the IBM Models 1-5, presented by Brown et al. (1993), and there are several implementations of these models, e.g., Giza++ by Och and Ney (2000), fast\_align by Dyer et al. (2013) or eflomal by Östling and Tiedemann (2016).

Word alignments used to be an integral part of statistical machine translation (SMT), and although they are not necessary for neural machine translation (NMT), they can be incorporated into different techniques to improve NMT (e.g., Alkhouli et al. (2018) which use explicit alignments for their NMT system). They are also still in use in other subfields of NLP, such as improving bilingual lexicon induction (Shi et al., 2021) or filtering noisy parallel corpora (Kurfali and Östling, 2019). See Steingrímsson et al. (2021) for a short overview.

More recently, word alignment systems that take advantage of contextualized word embeddings that are extracted from MLMs have been proposed. These systems compute word alignments by extracting from an MLM the embeddings of words in a sentence pair and putting them in a similarity matrix. The system then computes for each em-

bedding vector representing a word in the source sentence the most similar vector or vectors representing the words in the target sentence, which are then aligned. The major advantage of similarity-based word alignment is that the alignment quality is not affected by the amount of parallel data. Therefore, high-quality word alignments can be computed also in low-resource settings, when parallel data is scarce (Jalili Sabet et al., 2020). Two systems that use this concept are SimAlign by Jalili Sabet et al. (2020) and awesome-align by Dou and Neubig (2021).

## 4 Gold Standard for Word Alignment

In order to evaluate different word-alignment systems, a gold standard is needed (Koehn, 2009, p. 115). To this end, we created a gold standard consisting of 600 German-Romansh sentence pairs annotated for word alignment. The annotations were done using AlignMan, an annotation tool created by Steingrímsson et al. (2021).

### 4.1 Guidelines

Since descriptions of gold standard annotations for word alignment are rare in the literature, we decided to elaborate here on the annotation process.

To ensure consistency, a set of annotation guidelines were formulated. Following previous work, we aimed to “[a]lign as small segments as possible and as long segments as necessary.” (Ahrenberg (2007) citing Véronis and Langlais (2000)). More explicitly, we adhered to three general principles:

#### Principle I. Align only unambiguous words.

Since the annotation tool we used (ManAlign by Steingrímsson et al. (2021)) does not allow the use of confidence labels, i.e., it is only possible to annotate Sure but not Possible alignments, we only aligned words which are unambiguously mutual translations.

#### Principle II. Prefer 1-to-1 alignments over 1-to-n alignments or n-to-n alignments.

Since all alignments are Sure alignments, and since we evaluate word alignment, and not phrase alignment systems, 1-to-n alignments should be avoided. However, if a single word in the source sentence lexically corresponds to several words in the target sentence, they are aligned. In case words are repeated in one language but not in the other language, they are aligned only once.

<sup>5</sup>The dictionary was downloaded from <https://www.pledarigrond.ch/rumantschgrischun>

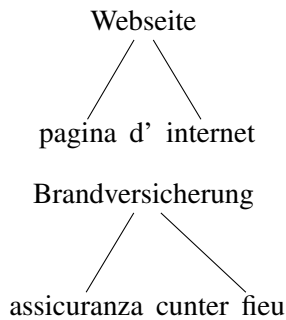


Figure 1: Aligning the German compounds *Webseite* (“website”) and *Brandversicherung* (“fire insurance”) to Romansh noun phrases. Only lexemes are aligned with each other. Romansh Prepositions are left unaligned.

**Principle III. Prefer lexical alignment over other types of alignments such as part-of-speech alignments or morphosyntactical alignments.** Only words that are translations of each other also outside the specific context of the sentence pair should be aligned. In cases of paraphrasing during translations, words should remain unaligned.

## 4.2 Annotation Challenges

We will now give some examples of the challenges of aligning German and Romansh words.

### 4.2.1 Compound Words

One difficulty is the alignment of German compounds (formations of new lexemes by adjoining two or more lexemes (Bauer, 1988)), to their Romansh translations, which are single words connected by a preposition (usually *da*). We decided to align German compounds to their corresponding Romansh lexemes, leaving the Romansh prepositions unaligned. See Figure 1 for two alignment examples.

### 4.2.2 German Preterite vs. Romansh Perfect

In the corpus at hand, two tenses are used in German for referring to past events: the preterite and the perfect. The German preterite is a synthetic verb form, i.e., it is made up of a single conjugated form, e.g., *nahm* (infinitive *nehmen* “take”) or *wurde* (infinitive *werden* “become”). The German perfect is an analytic construction, i.e., it is made up of an auxiliary verb (*haben* “have” or *sein* “be”) and the past participle, e.g., *Die Präsidentenkonferenz hat nun entschieden* (“the presidential conference has decided”).

Contrary to German, Romansh only has one tense referring to past events: the perfect. It is




Figure 2: Alignment of German preterite to Romansh perfect. The German word *war* is translated to Romansh *è stà*. Nonetheless, *è* is left unaligned since it only carries grammatical information (tense, number), but no lexical information.




Figure 3: The German verb *zurückweisen* (“reject, decline”), here separated into two words since it is used as the finite verb in the main clause, corresponds to the Romansh verb *renviar*. This results in a 2-to-1 alignment.

an analytic construction made of, similarly to German, an auxiliary *habere* “have” for transitive verbs or *esse* “be” for intransitive verbs and the past participle (Bossong, 1998, p. 189).

The German preterite (one word) is always translated using the Romansh perfect (two words). For example, in the sentence *Der Kanton Graubünden war letztmals 2003 Gastkanton* (“The last time the Canton of Grisons was a host canton was in 2003”) the verb *war* “was” is translated as *è stà*. This principally results in a 1-to-2 link. However, since in this context Romansh *è* only carries grammatical information of tense and number, but no additional lexical information, it remains unaligned, cf. Figure 2.

### 4.2.3 Separable Verbs

Separable verbs (German “*trennbare Verben*”) are verbs in front of which affixes (mostly prepositions) are placed. These affixes delimit and modify the verb’s meaning. In main clauses, if the finite verb is a separable verb, its affix will be placed at the end of the sentence, thus “seperable”. Since both the verb and the affix form together the lexical meaning of the word and are conceptually inseparable, both of them should be aligned to the corresponding Romansh verb, resulting in a 2-to-1 alignment. See Figure 3 for an example.

## 4.3 Shortcomings

Our gold standard has two shortcomings: First, it was created by a single annotator, the author of this

paper. Second, it only contains Sure alignments, which might influence the evaluation (cf. Section 5.1). The latter is however precedented, and there exist other word alignment gold standards which contain only Sure alignments, cf. Clematide et al. (2018) and Mihalcea and Pedersen (2003).

Refer to Dolev (2022, pp. 43-50) for further details about the creation of the gold standard as well as more annotation examples.

## 5 Experiments and Results

To test our hypothesis of whether MLMs contain information that can be meaningfully applied to Romansh, we compared two statistical word alignment systems as our baseline systems with two similarity- and embedding-based word alignment systems. For our baseline systems, we chose fast\_align (Dyer et al., 2013) and eflomal (Östling and Tiedemann, 2016). For our similarity-based systems, we chose SimAlign (Jalili Sabet et al., 2020) and awesome-align (Dou and Neubig, 2021).

Additionally, we tested the influence of the dataset size on the performance of the statistical models. We also experimented with fine-tuning mBERT using awesome-align on our parallel corpus.

### 5.1 Evaluation Metrics

To evaluate the models, we consider three measurements: precision, recall and alignment error rate (AER), which are defined as follows:

$$\text{Recall} = \frac{|A \cap S|}{|S|}, \quad \text{Precision} = \frac{|A \cap P|}{|A|}$$

$$\text{AER} = 1 - \frac{|A \cap S| + |A \cap P|}{|A| + |S|}$$

With  $A$  being the set of alignments generated by the model,  $S$  being the set of Sure alignments and  $P$  the set of Possible alignments, whereas  $S \subseteq P$ , meaning the set of Possible alignments also contains the Sure alignments (Och and Ney, 2000). In our case,  $P$  equals  $S$ , since our gold standard only has Sure alignments, hence  $P$  also stands for  $S$ .

To compute the measurements, we use the implementation from Jalili Sabet et al. (2020)<sup>6</sup>.




Figure 4: Performance of our baseline statistical models with relation to the dataset size.

### 5.2 Statistical Models

We tested the influence of dataset sizes on the statistical models, beginning with the 600 sentences from the gold standard, and up to the entire dataset (106,091 sentences). For that, we concatenated the 600 sentence pairs from the gold standard to parts of the dataset in different sizes, trained the models, and then took the alignments for the first 600 sentences, which correspond to the annotated sentence pairs. We then compared them to the gold alignments.

Surprisingly, fast\_align fails to match the performance that is reached by eflomal after just 600 sentences. As expected, the performance of both models increases with larger dataset sizes. Figure 4 visualizes the changes in precision and recall with dataset sizes. Figure 5 visualizes the AER of the baseline systems compared with the similarity-based systems. The exact precision, recall and AER values are given in Table 4 in Appendix A.

Putting these results into perspective, Östling and Tiedemann (2016) report an AER of 0.106 for English-Swedish (692,662 sentence pairs) and an AER of 0.279 for English-Romanian (48,641) sentence pairs (cf. Table 2 in Östling and Tiedemann (2016)). The performance of eflomal on our German-Romansh dataset, which reaches an AER of 0.146, is within this range.

### 5.3 Similarity-Based Models

To test whether MLMs contain information that can be meaningfully applied to Romansh, we compared word alignments computed by two similarity-based alignment models, which leverage contextualized word embeddings from MLMs. Since the

<sup>6</sup>[https://github.com/cisnlp/simalign/blob/master/scripts/calc\\_align\\_score.py](https://github.com/cisnlp/simalign/blob/master/scripts/calc_align_score.py)




Figure 5: Comparison of AER between the three systems (lower is better). The performance of fast\_align and eflomal profits from more data. The performance of SimAlign and awesome-align is not dependent on dataset size.

dataset size does not influence the performance of similarity-based word alignment systems, we only had the 600 sentence pairs from the gold standard aligned.

We tested the different algorithms (dubbed Argmax, Itermax and Match) offered by SimAlign to align words on a word and sub-word level and using both MLMs that are supported out-of-the-box by SimAlign: mBERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020). The embeddings were taken from the 8th layer, SimAlign’s default setting.

The results show that mBERT performs better than XLM-R, and that the Match algorithm performs best (Figure 6; The exact precision, recall and AER values are given in Table 5 in Appendix A).

As can be seen in Figure 5 and Table 2, both SimAlign and awesome-align (before fine-tuning) outperform fast\_align, but are still outperformed by eflomal.

To again put things into perspective, SimAlign’s performance on German-Romansh sentence pairs is on par with its performance on other language pairs: For the language pairs English-German, Jalili Sabet et al. (2020) report an AER of 0.19 using mBERT, which is only three percentage points better than its performance on German-Romansh. Also using mBERT, for the language pairs English-Hindi and English-Romanian, SimAlign achieves an AER of 0.45 and 0.35, respectively, which is considerably worse than its performance on German-Romansh. To conclude, SimAlign’s performance is within the range of what is reported by Jalili Sabet et al. (2020) for other language pairs. awesome-align

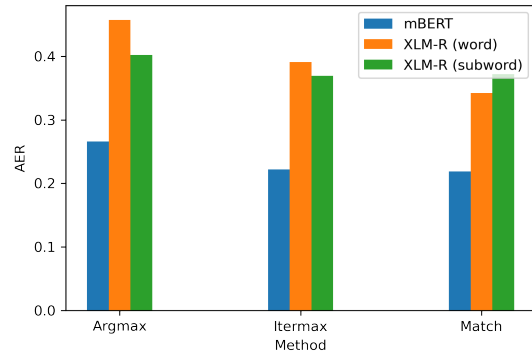


Figure 6: SimAlign’s AER using different language models and settings. Argmax, Itermax and Match refer to SimAlign’s algorithms for computing word alignment.

performs precisely as well as SimAlign.

### 5.3.1 Fine-tuning with parallel data

awesome-align offers the option of fine-tuning an MLM on a parallel corpus for improving word alignment. Dou and Neubig (2021)’s fine-tuning’s main objective consists of several sub-objectives. One sub-objective is simply continuing pre-training the language model in a masked language modeling objective on both sides of the parallel corpus. Another sub-objective, dubbed translation language modeling (TLM), consists of concatenating the parallel sentences and performing masked language modeling on the concatenated data. This objective “enable (*sic*) the model to align the source and the target representations.” (Dou and Neubig, 2021) Yet another sub-objective of fine-tuning, dubbed self-training objective (SO), is similar to the EM algorithm used in the IBM models. This objective brings the contextualized representations of words aligned in the first pass closer together, it reduces wrong alignments and encourages source-to-target and target-to-source alignments to be more symmetrical.

We fine-tuned mBERT using our German-Romansh parallel sentences to test the influence of fine-tuning using the settings recommended by Dou and Neubig (2021) on their GitHub page<sup>7</sup> which reportedly balance efficiency and effectiveness by applying TLM and SO.

With the fine-tuned model, awesome-align outperformed both baseline models and SimAlign, and reached an AER of 0.09. This performance lies

<sup>7</sup><https://github.com/neulab/awesome-align>

Method	Precision	Recall	AER
fast_align	0.625	0.787	0.307
eflomal	0.831	0.879	0.146
SimAlign	0.794	0.766	0.220
awesome-align	0.802	0.760	0.220
awesome-align (fine-tuned)	<b>0.912</b>	<b>0.911</b>	<b>0.090</b>

Table 2: Comparison of the best results of the baseline statistical models (above) with the best results of the two embedding-based models (below). Best result in bold.

between the AERs reported by [Dou and Neubig \(2021\)](#) for German-English (0.161) and French-English (0.041) after bilingual fine-tuning.

Based on this big performance improvement we hypothesize that by using parallel bilingual data for fine-tuning an MLM, a previously unseen language may be easily learned by the model. Thus, MLMs could be extended to include low-resource unseen languages, by focusing on parallel data with a seen language.

#### 5.4 Qualitative Evaluation

For a short qualitative evaluation of the different models, please refer to Appendix B.

#### 5.5 Conclusion

Both our conditions from Section 1 (similarity-based alignments should be on par with statistical models, and they should be on par with performance on seen languages) were fulfilled. SimAlign’s and awesome-align’s performance for German-Romansh sentence pairs (1) lies between that of our baseline models and (2) lies within the performance range on seen languages.

We see this as evidence supporting our hypothesis that MLMs, especially mBERT, contain information that can be meaningfully applied for processing Romansh.

We also showed that mBERT reacts well to fine-tuning on German-Romansh parallel data: after fine-tuning, the AER dropped from 0.22 to 0.09.

We thus believe that leveraging MLMs, especially mBERT, would be a promising way for developing language technology and NLP applications for Romansh.

## 6 Discussion

Multilingual models such as mBERT show good performance in a so-called “cross-lingual zero-shot transfer”: fine-tuning a pre-trained model on a task, e.g., PoS tagging, on one language and then carrying out this task on a different language (target language) for which it wasn’t trained ([Deshpande et al., 2022](#)). Such models also perform well in a variety of tasks such as PoS tagging or NER on unseen languages, i.e., languages which were not covered by the pre-trained model, such as Faroese, Maltese or Swiss German ([Muller et al., 2021](#)).

There is a lack of consensus as to what properties of a language favor performance in such scenarios, i.e., it is not entirely clear *when* zero-shot transfer works. Some suggest sub-word overlap is crucial for good performance ([Wu and Dredze, 2019](#)), while others show that transfer also works well between languages written in different scripts when they are typologically similar<sup>8</sup>, meaning sub-word overlap is not a necessary condition ([Pires et al., 2019](#)). It was, however, shown by [Muller et al. \(2021\)](#) that transliterating languages from unseen scripts leads to large gains in performance.

[Deshpande et al. \(2022\)](#) showed that zero-shot transfer is possible for different scripts with similar word order and that the lack of both, on the other hand, hurts performance. They also showed that zero-shot performance is correlated with alignment between word embeddings, i.e., to what extent the embeddings of different languages share the same geometric shape and are aligned across the same axes. [Muller et al. \(2021\)](#) ascribe differences in mBERT’s performance on unseen languages to two things: close relatedness to languages used during pretraining; and the unseen languages using the same script as those closely related languages that were seen during pretraining.

Romansh shares a high similarity, not only in script but also typologically, with other Romance languages, as well as other European languages<sup>9</sup>, which are a major part of the training data for mBERT. Therefore, we interpret the good performance of SimAlign and awesome-align as a reassurance of the assumptions cited above, namely that typological similarity and sub-word overlap favor zero-shot performance.

<sup>8</sup>mBERT fine-tuned for PoS tagging in Urdu (Arabic script) achieved 91% accuracy on Hindi (Devanagari script) ([Pires et al., 2019](#)). Both languages are mutually intelligible and are considered variants of a single language—Hindustani.

<sup>9</sup>European languages from different language families

## 7 Concluding Words

We presented a new trilingual corpus (German, Romansh, Italian), which is a continuation of the ALLEGRA corpus by Scherrer and Cartoni (2012). We also presented a gold standard for German-Romansh word alignments, consisting of 600 sentences.

We showed that word alignments computed using similarity-based word alignment models with embeddings taken from MLMs, specifically mBERT, are on par with statistical models; they are also on par with similarity-based models' performance on seen languages. We view this as a good indication that these models contain information that is meaningful and applicable to Romansh.

## 8 Ethical Considerations

The corpus we collected consists of press releases that are publicly available on the Internet. The copyright lies with the Canton of Graubünden. Downloading the press releases is explicitly allowed.

The State Chancellery of Grisons gave their approval that the data collected for this work and published therein be further available for research purposes. The commercial use of these documents is however prohibited. In any event, please refer to the copyright notice<sup>10</sup> on the canton's website.

### 8.1 Environmental Impact

Four hours were needed for fine-tuning mBERT for awesome-align on our GPU (nVIDIA GTX 1080). Total emissions are estimated at 0.31 kgCO<sub>2</sub>eq<sup>11</sup>. Estimations were conducted using the Machine-Learning Impact calculator presented in Lacoste et al. (2019). 100% of the electricity used for running our machine came from renewable energy.

## Acknowledgements

Eyal Liron Dolev is funded by the Swiss National Science Foundation (Sinergia project grant no. 205913). I would like to thank Prof. Martin Volk for supervising my master's thesis (cf. Dolev (2022)) which served as the basis for this paper. I

(Germanic, Romance, Slavic) were shown to display high similarity to each other and to form a so-called *Sprachbund*, dubbed Standard Average European (Haspelmath, 2001).

<sup>10</sup><https://www.gr.ch/de/Seiten/Impressum.aspx>

<sup>11</sup>The equivalent of 2.5 km driven by a European car, according to the European Parliament's website.

would also like to thank Jannis Vamvas for encouraging me to publish this work, as well as Tannon Kew for his most valuable feedback.

## References

- Lars Ahrenberg. 2007. *Lines 1.0 annotation: Format, contents and guidelines*. Technical report.
- Tamer Alkhouli, Gabriel Bretschner, and Hermann Ney. 2018. *On the alignment problem in multi-head attention-based neural machine translation*. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 177–185, Brussels, Belgium. Association for Computational Linguistics.
- Mikel Artetxe and Holger Schwenk. 2019. *Massively Multilingual Sentence Embeddings for Zero-Shot Cross-Lingual Transfer and Beyond*. *Transactions of the Association for Computational Linguistics*, 7:597–610.
- Laurie Bauer. 1988. *Introducing Linguistic Morphology*. Edinburgh University Press.
- Reto Baumgartner, Martina Bachmann, Rolf Badat, Daniel Hegglin, Susanna Tron, Melanie Widmer, Nora Löttscher, Noëmi Aepli, Martin Cantieni, and Victoria Mosca. 2014. *Morphologieanalyse für Rumantsch Grischun*. University of Zurich.
- Steven Bird, Edward Loper, and Ewan Klein. 2009. *Natural Language Processing with Python*. O'Reilly Media Inc.
- Georg Bossong. 1998. *Die Romanischen Sprachen: Eine vergleichende Einführung*. Helmut Buske Verlag, Hamburg.
- Peter F. Brown, Stephen A. Della Pietra, Vincent J. Della Pietra, and Robert L. Mercer. 1993. *The mathematics of statistical machine translation: Parameter estimation*. *Computational Linguistics*, 19(2):263–311.
- Simon Clematide, Stéphanie Lehner, Johannes Graën, and Martin Volk. 2018. *A multilingual gold standard for translation spotting of German compounds and their corresponding multiword units in English, French, Italian and Spanish*, pages 125–145. John Benjamins.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. *Unsupervised cross-lingual representation learning at scale*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Ameet Deshpande, Partha Talukdar, and Karthik Narasimhan. 2022. *When is BERT multilingual? Isolating crucial ingredients for cross-lingual transfer*.



- In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3610–3623, Seattle, United States. 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.
- Eyal Liron Dolev. 2022. [Using multilingual word embeddings for similarity-based word alignments in a zero-shot setting: Tested on the case of German-Romansh](#). Master’s thesis, University of Zurich.
- Zi-Yi Dou and Graham Neubig. 2021. [Word alignment by fine-tuning embeddings on parallel corpora](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2112–2128, Online. Association for Computational Linguistics.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. [A simple, fast, and effective reparameterization of IBM model 2](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia. Association for Computational Linguistics.
- Johannes Graën, Dominique Sandoz, and Martin Volk. 2017. [Multilingwis2 – explore your parallel corpus](#). In *Proceedings of the 21st Nordic Conference on Computational Linguistics, NoDaLiDa*, number 131 in Linköping Electronic Conference Proceedings, pages 247–250. Linköping University Electronic Press, Linköpings universitet.
- John Haiman and Paola Benincà. 1992. *The Rhaeto-Romance Languages*. Routledge, London and New York.
- Martin Haspelmath. 2001. [The European linguistic area: Standard Average European](#). *Language Typology and Language Universals*, 20.2:1492–1510.
- Masoud Jalili Sabet, Philipp Dufter, François Yvon, and Hinrich Schütze. 2020. [SimAlign: High quality word alignments without parallel training data using static and contextualized embeddings](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1627–1643, Online. Association for Computational Linguistics.
- Philipp Koehn. 2009. *Statistical Machine Translation*. Cambridge University Press.
- Murathan Kurfalı and Robert Östling. 2019. [Noisy parallel corpus filtering through projected word embeddings](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 277–281, Florence, Italy. Association for Computational Linguistics.
- Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. [Quantifying the carbon emissions of machine learning](#). *arXiv preprint arXiv:1910.09700*.
- Rada Mihalcea and Ted Pedersen. 2003. [An evaluation exercise for word alignment](#). In *Proceedings of the HLT-NAACL 2003 Workshop on Building and Using Parallel Texts: Data Driven Machine Translation and Beyond*, pages 1–10.
- Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot, and Djamé Seddah. 2021. [When being unseen from mBERT is just the beginning: Handling new languages with multilingual language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 448–462, Online. Association for Computational Linguistics.
- Mathias Müller, Annette Rios, and Rico Sennrich. 2020. [Domain robustness in neural machine translation](#). In *Proceedings of the 14th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 151–164, Virtual. Association for Machine Translation in the Americas.
- Franz Josef Och and Hermann Ney. 2000. [Improved statistical alignment models](#). In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, pages 440–447, Hong Kong. Association for Computational Linguistics.
- Robert Östling and Jörg Tiedemann. 2016. [Efficient word alignment with Markov Chain Monte Carlo](#). *Prague Bulletin of Mathematical Linguistics*, 106:125–146.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Yves Scherrer and Bruno Cartoni. 2012. [The trilingual ALLEGRA corpus: Presentation and possible use for lexicon induction](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2890–2896, Istanbul, Turkey. European Language Resources Association (ELRA).
- Haoyue Shi, Luke Zettlemoyer, and Sida I. Wang. 2021. [Bilingual lexicon induction via unsupervised bitext construction and word alignment](#). 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 813–826, Online. Association for Computational Linguistics.

Steinþór Steingrímsson, Hrafn Loftsson, and Andy Way. 2021. [CombAlign: a tool for obtaining high-quality word alignments](#). In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 64–73, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.

Dániel Varga, László Németh, Péter Halácsy, András Kornai, Viktor Trón, and Viktor Nagy. 2005. [Parallel corpora for medium density languages](#). In *Proceedings of the RANLP 2005*, pages 590–596.

Jean Véronis and Philippe Langlais. 2000. [Evaluation of parallel text alignment systems. The ARCADE project](#), pages 369–388. Springer Netherlands, Dordrecht.

Martin Volk and Simon Clematide. 2014. [Detecting code-switching in a multilingual alpine heritage corpus](#). In *Proceedings of the First Workshop on Computational Approaches to Code Switching*, pages 24–33, Doha, Qatar. Association for Computational Linguistics.

Manuela Weibel. 2014. [Aufbau paralleler Korpora und Implementierung eines wortalignierten Suchsystems für Deutsch–Rumantsch Grischun](#). Master’s thesis, University of Zurich.

Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 833–844, Hong Kong, China. Association for Computational Linguistics.

## A Tables

Table 3 gives information about the number of documents in our corpus for each language and year, as well as the number of documents that exist in at least two languages. Table 4 provides precision, recall and AER measurements for the statistical baseline models, tested on different dataset sizes (see also Figures 4 and 5). Table 5 gives information about precision, recall and AER for SimAlign’s different algorithms, using different embeddings and tokenization (see also Figure 6).

## B Qualitative Evaluation

We attempt to give a short qualitative evaluation of the advantages and shortcomings of the models we evaluated, i.e., describe which linguistic features the models deal with well and which features make them fail. To this end, we chose three sample sentences containing two potential pitfalls: aligning compound words and aligning differing word order. We compare our annotations, i.e., the gold

Year	DE	RM	IT	Parallel
1997	181	17	18	11
1998	168	153	153	147
1999	161	130	130	131
2000	192	167	169	168
2001	233	159	171	172
2002	235	157	165	153
2003	167	110	111	102
2004	132	97	94	89
2005	157	134	133	127
2006	211	173	174	166
2007	199	147	145	140
2008	201	168	169	164
2009	212	175	176	109
2010	219	183	184	184
2011	203	167	167	167
2012	254	207	207	207
2013	260	219	219	219
2014	260	218	218	218
2015	227	183	183	183
2016	221	190	190	190
2017	236	207	207	207
2018	248	221	220	221
2019	238	216	216	216
2020	310	284	285	286
2021	322	294	294	294
2022	301	275	276	276
Total	5,748	4,651	4,674	4,547

Table 3: Number of documents in each language per year. “Parallel” is the number of documents that exist in at least two languages.

Method	Dataset Size	Precision	Recall	AER
fast_align	600	0.515	0.644	0.427
	1k	0.526	0.662	0.414
	5k	0.561	0.704	0.375
	10k	0.579	0.725	0.356
	25k	0.601	0.75	0.332
	50k	0.615	0.768	0.317
	75k	0.626	0.781	0.305
	100k	<b>0.625</b>	<b>0.787</b>	<b>0.307</b>
eflomal	600	0.707	0.724	0.284
	1k	0.732	0.742	0.263
	5k	0.776	0.78	0.222
	10k	0.798	0.805	0.199
	25k	0.812	0.836	0.176
	50k	0.828	0.85	0.161
	75k	0.827	0.856	0.159
	100k	<b>0.831</b>	<b>0.879</b>	<b>0.146</b>

Table 4: Word alignment quality of the baseline statistical models, tested on different dataset sizes. Best result per method in bold. “Dataset Size” refers to the number of sentence pairs. The full dataset size is the number of sentence pairs extracted at the time of the experiments (106,691) plus the 600 annotated sentence pairs from the gold standard.

standard, with eflomal, the best statistical model, with SimAlign (using embeddings from mBERT and the Match method) and with awesome-align, before and after fine-tuning.

In the following plots, the dark-colored squares (green/blue) represent the gold alignment. The green plots compare eflomal and SimAlign. The blue plots compare awesome-align before and after fine-tuning. The plots were created using a slightly altered version of a script made available by Jalili Sabet et al. (2020) on their GitHub page<sup>12</sup>.

It seems that of all systems, SimAlign has the biggest difficulty aligning German compounds to their Romansh counterparts. In Figure 7, SimAlign (green, boxes), fails to align the words *Medienpaket* and *Medienvielfalt*, whereas eflomal (green, circles) manages to align at least *Medienvielfalt* successfully. awesome-align (blue, circles) manages to align both compounds correctly but still has some precision problems. After fine-tuning (blue, boxes), awesome-align’s accuracy is perfect.

The sentence in Figure 8 is an example of differing word order between German and Romansh. All models seem to have difficulties aligning this sentence. eflomal (green, circles) performs best with

an accuracy of 0.66, followed by awesome-align (blue, circles) and SimAlign (green, boxes), with an accuracy of 0.61 and 0.58, respectively. After fine-tuning, awesome-align (blue, boxes), outperforms the other models with an accuracy of 0.92.

The sentence in Figure 9 is a general example of an unsuccessful alignment: Both SimAlign and awesome-align fail, although it is not quite clear why (presumably too many Romansh words are not similar enough to words in seen languages). SimAlign (green, boxes) and awesome-align (blue, circles) have a very low accuracy of 0.33 and 0.4, respectively. eflomal, the statistical model (green, circles), reaches an accuracy of 0.86. However, after fine-tuning, awesome-align (blue, boxes) seems to have learned the right word correspondences and reaches perfect accuracy.

To conclude, a superficial and unsystematic comparison between the systems gives the impression that both SimAlign and awesome-align have difficulties aligning compounds and differing word order. Romansh words that are less similar to words in seen languages might also have a negative impact on this. As the case may be, fine-tuning mBERT using awesome-align’s training objectives helps awesome-align deal with those difficulties and shows a big improvement in accuracy.

<sup>12</sup><https://github.com/cisnlp/simalign/blob/master/scripts/visualize.py>




Figure 7: Example for aligning compounds. Green: eflomal (boxes) vs. SimAlign (circles); Blue: awesome-align (circles) vs. after fine-tuning (boxes). The dark green/blue squares represent the gold alignments.

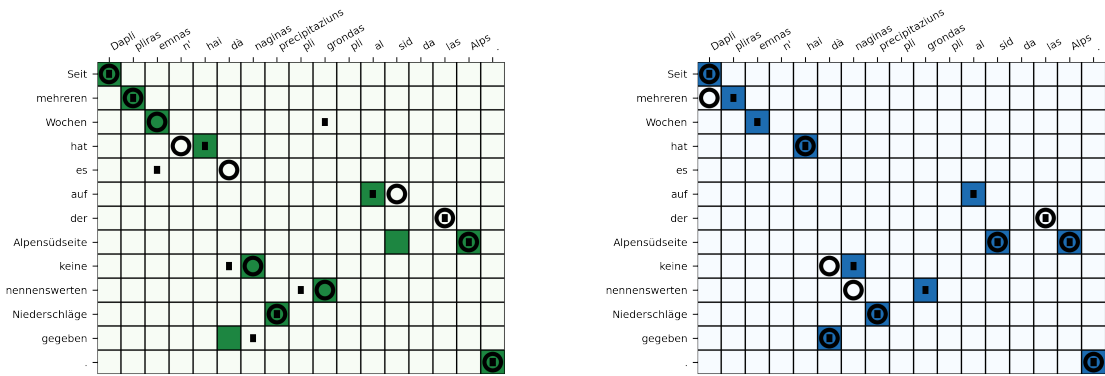


Figure 8: Example for aligning a sentence with differing word order. Green: eflomal (boxes) vs. SimAlign (circles); Blue: awesome-align (circles) vs. after fine-tuning (boxes). The dark green/blue squares represent the gold alignments.

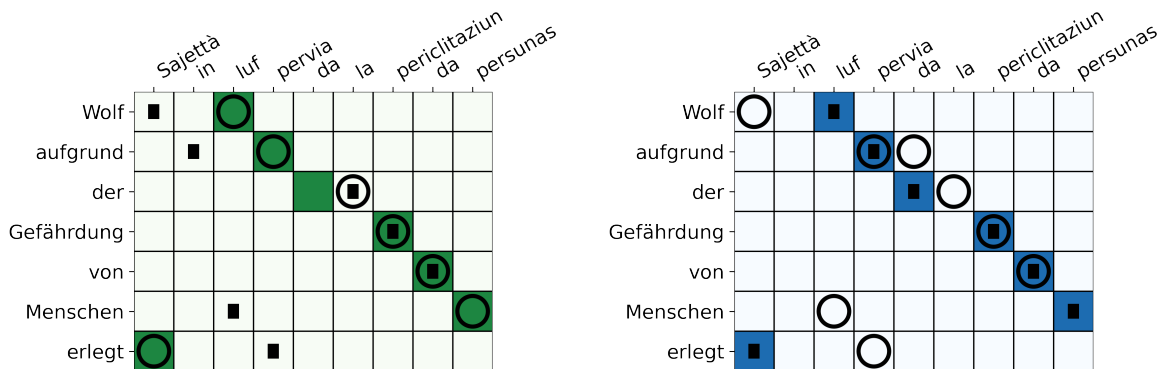


Figure 9: Example for the clear improvement after fine-tuning. Green: eflomal (boxes) vs. SimAlign (circles); Blue: awesome-align (circles) vs. after fine-tuning (boxes). The dark green/blue squares represent the gold alignments.

	Model	Level	Method	Precision	Recall	AER
SimAlign	mBERT	Subword	Argmax	<b>0.894</b>	0.622	0.266
			Itermax	0.832	0.731	0.222
			Match	0.795	<b>0.781</b>	<b>0.220</b>
	XLM-R	Word	Argmax	<b>0.848</b>	0.399	0.457
			Itermax	0.767	0.504	0.391
			Match	0.67	0.647	<b>0.342</b>
	XLM-R	Subword	Argmax	0.773	0.488	0.402
			Itermax	0.671	0.595	0.369
			Match	0.558	<b>0.719</b>	0.372

Table 5: Word alignment quality using SimAlign, with embeddings taken from different MLMs using word/sub-word level. Embeddings were taken from the eighth layer (default SimAlign setting). Best result per embedding type in bold.