



ACE-2005-PT: Corpus for Event Extraction in Portuguese

Luís Filipe Cunha

lfc@uminho.pt

FCUP-University of Porto

INESC Technology and Science

Porto, Portugal

University of Minho - Braga, Portugal

Ricardo Campos

ricardo.campos@ubi.pt

University of Beira Interior

Covilhã, Portugal

INESC Technology and Science

Porto, Portugal

Ci2 - Smart Cities Research Center - IP Tomar

Tomar, Portugal

Purificação Silvano

msilvano@letras.up.pt

FLUP-University of Porto,

CLUP

Porto, Portugal

Alípio Jorge

amjorge@fc.up.pt

FCUP-University of Porto

INESC Technology and Science

Porto, Portugal

ABSTRACT

Event extraction is an NLP task that commonly involves identifying the central word (trigger) for an event and its associated arguments in text. ACE-2005 is widely recognised as the standard corpus in this field. While other corpora, like PropBank, primarily focus on annotating predicate-argument structure, ACE-2005 provides comprehensive information about the overall event structure and semantics. However, its limited language coverage restricts its usability. This paper introduces ACE-2005-PT, a corpus created by translating ACE-2005 into Portuguese, with European and Brazilian variants. To speed up the process of obtaining ACE-2005-PT, we rely on automatic translators. This, however, poses some challenges related to automatically identifying the correct alignments between multi-word annotations in the original text and in the corresponding translated sentence. To achieve this, we developed an alignment pipeline that incorporates several alignment techniques: lemmatization, fuzzy matching, synonym matching, multiple translations and a BERT-based word aligner. To measure the alignment effectiveness, a subset of annotations from the ACE-2005-PT corpus was manually aligned by a linguist expert. This subset was then compared against our pipeline results which achieved exact and relaxed match scores of 70.55% and 87.55% respectively. As a result, we successfully generated a Portuguese version of the ACE-2005 corpus, which has been accepted for publication by LDC.

CCS CONCEPTS

• **Computing methodologies** → *Language resources; Information extraction; Machine translation.*

KEYWORDS

Corpus Translation, Annotation Alignment, Event Extraction

ACM Reference Format:

Luís Filipe Cunha, Purificação Silvano, Ricardo Campos, and Alípio Jorge. 2024. ACE-2005-PT: Corpus for Event Extraction in Portuguese. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, July 14–18, 2024, Washington, DC, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3626772.3657872>

1 INTRODUCTION

Event extraction is a crucial Information Extraction task, that has witnessed a growing interest among researchers in recent years [14]. However, this task can pose significant challenges due to its inherent complexity and the diverse ways in which events can be expressed in natural language [13]. Despite these challenges, several works [2, 16–18] have demonstrated significant progress in this field, achieving promising results.

One of the most widely used corpora in event extraction research is the ACE-2005 (Automatic Content Extraction) corpus [9], which has become a gold standard in the field [7]. This corpus provides a comprehensive collection of annotated events, offering a valuable benchmark for evaluating the effectiveness of event extraction systems. Unfortunately, it is only available in English, Chinese and Arabic, but not in other languages, such as Portuguese.

In order to bridge this gap and promote event extraction for the Portuguese language, we introduce ACE-2005-PT (Automatic Content Extraction for Portuguese), an automated translation of the ACE-2005 corpus that has been partially verified by a linguist expert. To generate this corpus, we developed a translation and annotation alignment pipeline that combines machine translation and several text alignment techniques such as fuzzy string matching, lemmatization, synonym matching, multiple translations and a word aligner based on Transformers [21]. By leveraging this pipeline, we have successfully translated the ACE-2005 corpus into Portuguese, including two variants: European Portuguese and Brazilian Portuguese. Then, a subset of ACE-2005-PT annotations



This work is licensed under a Creative Commons Attribution International 4.0 License.

SIGIR '24, July 14–18, 2024, Washington, DC, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0431-4/24/07

<https://doi.org/10.1145/3626772.3657872>

was manually aligned by a linguist to validate the pipeline’s effectiveness.

This corpus has already been accepted for publication by the Linguistic Data Consortium (LDC) and will be available for event extraction-related tasks. The main contributions of this paper are:

- An automatic translation and alignment pipeline that can be used to translate not only ACE-2005 but other corpora into different languages. We made the code available¹ for reproducibility purposes.
- A Portuguese-translated version of the ACE-2005 corpus that can be used to enhance Event Extraction in two Portuguese variants: European and Brazilian.

In the following sections, we present a brief description of ACE-2005 and the methodology used to translate it into Portuguese.

2 ACE-2005 CORPUS

The ACE-2005 corpus is composed of several textual documents from various sources, such as newswires, online journals, broadcast transcripts, discussion forums and conversational telephone speech. Each document is provided with annotations of events, which consist of event triggers and their corresponding event arguments [1]. Event triggers represent the terms that indicate the occurrence of an event and play a crucial role in event extraction. Each event trigger is associated with a specific event type resulting in 33 event types in the ACE-2005 corpus. Event arguments describe entities, temporal expressions or values that serve as participants or attributes of the events. Each event argument is associated with a semantic role that represents its relationship within the event, such as the agent that performs an action and the time or location of the event. For each event occurrence, ACE-2005 annotates event arguments by linking them to their respective event trigger.

Consider the following sentence, which illustrates the process of extracting an event trigger and its corresponding arguments.

Marie Curie was born in Warsaw on November 7, 1867.

In this example, the word "born" corresponds to a trigger with type Life:Be-Born. Regarding the event arguments, we have "Marie Curie", "Warsaw" and "November 7, 1867" with roles Person, Place and Time respectively.

3 ACE-2005 TRANSLATION TO PORTUGUESE

In this section, we introduce the translation process and annotation alignment pipeline used to generate the ACE-2005-PT corpus and present statistics associated with this methodology. Following that, we introduce the methods used to evaluate our pipeline effectiveness and the obtained results.

3.1 Translation and Annotation Alignment

The initial step in the process of producing the ACE-2005 corpus consisted of performing automatic translation of the documents. Specifically, we used Google Translator² for Brazilian Portuguese, while for European Portuguese, we relied on DeepL translator³, as

Google Translator does not support translation into this language variant.

After translation, we have an original text with annotated terms and we want to transfer these annotations to the translated text, in a process called alignment (of annotations). In the simplest case, we translate the annotated term and its translation becomes annotated in the translated text. However, such a direct approach works only half of the time. For instance, in the sentence “The soldiers were ordered to fire their weapons”, ACE-2005 states that the trigger “fire” should be annotated. However, this sentence is translated to “Os soldados receberam ordens para disparar as suas armas” where the word “fire” is translated to “incêndio” (fire as a noun) in isolation and to “disparar” (fire as a verb) in context.

We employed a commonly used pre-processing approach⁴ of the ACE-2005 corpus, which performed sentence tokenization on the documents and assigned each event annotation to its respective sentence. Subsequently, we automatically translated each source (*src*) sentence s_{src} , along with its corresponding triggers t_{src} and arguments a_{src} . In case a mismatch occurs, these translations (*trans*) result in annotations a_{trans} and t_{trans} that are not contained in the sentences s_{trans} . For these cases, we applied an alignment procedure (Algorithm 1) to identify the correct span offsets of a_{trans} and t_{trans} within s_{trans} . For the sake of simplicity, we only describe the event argument alignment process as the trigger alignment procedure was similar. The alignment pipeline is composed of four components: lemmatization, multiple translations, a BERT-based word aligner and fuzzy string similarity.

Algorithm 1: Translation and Alignment Pipeline.

```

1 for  $s_{src}$  in sentences do
2    $s_{trans} \leftarrow$  get sentence translation;
3   for  $a_{src}$  in event arguments do
4      $a_{trans} \leftarrow$  get argument translation;
5     if  $a_{trans}$  in  $s_{trans}$  then return  $a_{trans}$  ;
6      $match \leftarrow$  LemmaMatch( $a_{lem}, s_{lem}$ );
7     if  $match$  then return  $match$  ;
8      $match \leftarrow$  MultipleTranslation( $a_{src}, s_{trans}$ );
9     if  $match$  then return  $match$  ;
10     $match \leftarrow$  WordAligner( $s_{src}, s_{trans}, a_{src}$ );
11    if  $match$  then return  $match$  ;
12    return FuzzyMatch( $a_{trans}, s_{trans}$ );
13  end
14 end
```

The initial step of the alignment pipeline starts by lemmatizing both s_{trans} and a_{trans} , resulting in s_{trans}^{lem} and a_{trans}^{lem} , as shown in Algorithm 2. Next, we verify whether the sequence of tokens in a_{trans}^{lem} occurs in s_{trans}^{lem} . By transforming each token of s_{trans} and a_{trans} into their fundamental word forms, we reduce the potential variations for each word and increase the likelihood of finding a match.

The next element of the pipeline consists of using multiple translations of words corresponding to annotated terms. In particular,

¹<https://github.com/LIAAD/ACE-2005-Translation-and-Alignment-Pipeline>

²<https://cloud.google.com/translate/docs/reference/rest>

³<https://www.deepl.com>

⁴<https://github.com/nlpl-lab/ace2005-preprocessing>

Algorithm 2: LemmaMatch function.

```

1 def getLemmaMatch( $s_{trans}$ ,  $a_{trans}$ ):
2    $a_{trans}^{lem} \leftarrow$  get  $a_{trans}$  lemma tokens;
3    $s_{trans}^{lem} \leftarrow$  get  $s_{trans}$  lemma tokens;
4   if  $a_{trans}^{lem}$  in  $s_{trans}^{lem}$  then
5      $start, end \leftarrow$  find_span( $s_{trans}^{lem}$ ,  $a_{trans}^{lem}$ );
6     return  $a_{trans} \leftarrow s_{trans}[start, end]$ ;

```

we used the Microsoft Dictionary Lookup API⁵ which provides alternative translations for words and idiomatic expressions to generate alternative translations of a_{trans} and t_{trans} , aiming to locate them in s_{trans} . This method was also combined with the previous one by performing lemmatization on the alternative translation and trying to match it within the s_{trans} .

Following these methods, we used a parallel corpus word aligner adapted from the word aligner proposed by Dou et al. (2021) [10]. We retrieved the embedding matrices of s_{src} and s_{trans} from the BERT-Multilingual model (mBERT) [8] and computed their product. Subsequently, we applied a softmax function to compute probabilities for the association between tokens in s_{src} and s_{trans} .




Figure 1: Alignment matrix from mBERT.

Figure 1 presents an example of the embeddings-driven word association matrix with an English sentence s_{src} on the vertical axis and its corresponding Portuguese translation s_{trans} on the horizontal axis. With the word alignments between s_{src} and s_{trans} computed, we proceeded to find the correct a_{trans} span offset. For that, we used an approach represented in Algorithm 3, which is similar to the one described by Carrino et al. (2020)[4], with the inclusion of a token size verification. First, we calculate the start (a_{src}^{start}) and end (a_{src}^{end}) positions of a_{src} in s_{src} . Next, we use the

word alignments computed in $align(s_{src}, s_{trans})$ to retrieve the corresponding start and end positions of a_{trans} in s_{trans} , retrieving a_{trans}^{start} and a_{trans}^{end} respectively. Finally, we determine the corresponding a_{trans} by using its span offset as indexes in s_{trans} . Then we verify the length of the aligned annotation by comparing it with the original annotation length to avoid returning disproportionately large annotations. (The motivation for this is further explained in Section 4).

Algorithm 3: Word Aligner Algorithm

```

1 def WordAligner( $s_{src}$ ,  $s_{trans}$ ,  $a_{src}$ ):
2    $aligns \leftarrow$  align( $s_{src}$ ,  $s_{trans}$ ); /*mBERT*/
3   ( $a_{src}^{start}, \dots, a_{src}^{end}$ )  $\leftarrow$  get  $a_{src}$  span offsets;
4   ( $a_{trans}^{start}, \dots, a_{trans}^{end}$ )  $\leftarrow$   $a_{trans}$  from aligns;
5    $a_{trans}^{start} \leftarrow$  min( $a_{trans}^{start}, a_{trans}^{end}$ );
6    $a_{trans}^{end} \leftarrow$  max( $a_{trans}^{start}, a_{trans}^{end}$ );
7    $a_{trans} \leftarrow$   $s_{trans}[a_{trans}^{start}, a_{trans}^{end}]$ ;
8   if SizeSafeguard( $a_{trans}$ ,  $a_{src}$ ) then
9     return  $a_{trans}$ 
10  end

```

The last component of the pipeline uses Fuzzy string similarity. We calculate the similarity between a_{trans} and substrings of s_{trans} . To generate these substrings, we use a sliding window to compute all substrings of s_{trans} with length $window$ equal to the number of tokens in a_{trans} . In order to determine the most similar string between a_{trans} and the generated substrings, we use fuzzy matching algorithms such as the Levenshtein distance [12] and Gestalt pattern matching [20].

In the case of trigger alignment, where triggers generally consist of a single word, we did not use the fuzzy match component. This decision was based on the limitations associated with character-level similarity techniques, which are further elaborated in Section 4. Instead, we used synonyms of t_{trans} , trying to establish a correspondence between t_{trans} and s_{trans} .

3.2 Statistics by pipeline component

Using the alignment pipeline, we produced ACE-2005-PT, featuring 16,260 sentences and 14,886 annotations (triggers and arguments), mirroring the original corpus. In this section, we focus, for the sake of simplicity, on the European corpus, as the results were similar in the Brazilian variant. Upon translation, we were able to find 2,721 triggers (51.9%) and 5,127 arguments (53.1%) in their respective translated sentences by using substring matching between the annotations and the translated text. This indicates that approximately half of the ACE-2005-PT annotations required alignment.

Table 1 provides the number of alignments performed by each component of our alignment pipeline for each data split. Notably, the pipeline works sequentially, meaning that annotations aligned by earlier methods are not addressed again by subsequent pipeline elements. To select the best component order in the pipeline we experimented with all the permutations between the components and calculated the corresponding alignment results using a manually aligned corpus that is further introduced in Section 3.3. The

⁵<https://learn.microsoft.com/en-us/azure/ai-services/translator/reference/v3-0-dictionary-lookup>

Table 1: ACE-2005-PT alignment statistics. Number of alignments by component. (SMatch - string match; MTrans - multiple translation; WAligner - word aligner)

Method	Trigger				Argument			
	Train	Dev	Test	Total	Train	Dev	Test	Total
SMatch	2,289	222	210	2,721	4,109	503	515	5,127
Lemma	473	56	38	567	1,235	152	136	1,523
MTrans	947	144	125	1,216	568	53	52	673
Synonym	16	2	1	19	-	-	-	-
WAligner	587	66	47	700	1,675	198	168	2,041
Fuzzy	-	-	-	-	152	25	16	193
Manual	12	2	1	15	84	2	5	91
Total	4,324	492	422	5,238	7,823	933	892	9,648

obtained results confirmed that the order that presents the best results is the one presented in Table 1 (top to bottom).

Despite some entries couldn't be automatically aligned using our pipeline, their number was marginal compared to the total number of annotations. After using the alignment pipeline, we conducted a set of 106 manual alignments, corresponding to the "Manual" method in Table 1.

3.3 Evaluation

To measure the effectiveness of the alignment pipeline, manual alignments were conducted on the entire ACE-2005-PT test set, which includes 1,310 annotations (triggers and arguments). These alignments were performed by an expert linguist to ensure high-quality annotations, following the same annotation guidelines of the original ACE-2005 corpus [1]. These annotations were then compared with the automatic alignment performed by our pipeline. During this comparison, we used two different metrics: the exact match between the pipeline-generated alignments and the manual ones (strict), and the F1 score between the tokens found by the pipeline and the ones that were manually aligned (relaxed). Resorting to both evaluation metrics allow us to more accurately assess the model's effectiveness.

3.4 Results

Table 2 shows the alignment validation results of our pipeline when compared with the manually annotated test split of ACE-2005-PT.

In this table, we can analyse the relaxed and exact match between the pipeline-generated alignments and the gold annotations performed by the linguistic expert. The results are presented for each alignment method, along with the overall performance of the pipeline. The pipeline achieved a Relaxed score of 87.72% and an exact score of 70.55%. As expected, the flexible match obtained a higher result value than the exact match. Looking at the overall performance, there is a 17.22% difference between both metrics, which shows that, despite our pipeline getting most of the alignments in the correct form, it still presents difficulties in reproducing the exact alignments performed by an expert.

Looking at the annotation types (triggers and arguments), our pipeline demonstrates more difficulties when aligning arguments

Table 2: Evaluation Results by pipeline component.

Method	Trigger		Argument		All	
	Relaxed	Exact	Relaxed	Exact	Relaxed	Exact
SMatch	99.29	99.05	92.21	76.36	94.25	82.92
Lemma	93.86	92.11	74.06	20.74	78.41	36.42
MTrans	93.73	90.40	77.04	46.15	88.83	77.40
Synonym	1.0	1.0	-	-	1.0	1.0
WAligner	63.54	58.33	68.26	28.57	67.22	35.19
Fuzzy	-	-	92.86	56.25	92.86	56.25
Pipeline	93.11	91.23	85.16	60.76	87.77	70.55

(85.16% relaxed and 60.76% exact) compared to triggers (93.11% relaxed and 91.23% exact). This is expected since a trigger is usually composed of one word, while arguments may contain several words, making it harder to identify the correct alignment. Another observation to be made is that in the case of trigger alignment, the results obtained for both relaxed and exact scores are similar (93.11% relaxed and 91.23% exact). On the other hand, one can observe, in the argument alignments, a higher difference between the relaxed and the exact score, with a score of 85.15% and 60.76% respectively. Again, this indicates that aligning arguments is a harder task for our pipeline.

Finally, looking at the aligning modules individually, one can observe that the string match of the translated annotation with the translated text is the method that yielded the highest evaluation scores, for both triggers and arguments. Both the Lemmatization (Lemma) and the Multiple Word Translation (MTrans) methods revealed alignment scores above 90% on the trigger alignment. As for the argument alignment, the Fuzzy string match and MTrans were the ones that achieved better alignment scores. The limitations of each alignment technique are further explored in Section 4.

The quality of these alignments can further be assessed by observing the Event Extraction task results of models that were trained using the ACE-2005-PT corpus. In Cunha et. al (2023) [6] this corpus was used to train Question Answering models to perform Event Extraction for Portuguese. The model trained with ACE-2005-PT is available for direct scrutiny online⁶.

Table 3 presents F1-scores for trigger extraction and argument extraction on both the ACE-2005 corpus in English (original) and our Portuguese version presented in this paper, ACE-2005-PT.

Since the Portuguese model consists of a QA model it should primarily be compared with the BERT_QA_Arg model in Table 3 due to the similarities in their model architectures. Looking at the table, one can observe that the results for the Portuguese model (64% F1-score for Trigger Extraction and 46% F1-score for Argument Extraction) are lower when compared to the results of the English models. However, since the language used to train and evaluate these models is different, a direct comparison might be inaccurate due to divergences in language vocabulary, such as idiomatic expressions, and other cultural differences between languages.

⁶<https://hf.co/spaces/lfcc/Event-Extractor>

Table 3: Event Extraction results on ACE-2005 dataset.

Model	Triggers F1	Arguments F1
English ACE-2005		
JRNN 2016 [19]	69.3	55.4
BERT_QA_Arg 2020 [11]	72.4	53.3
OneIE 2020 [15]	74.7	56.8
Text2Event 2021 [16]	71.9	53.8
FourIE 2021 [17]	75.4	58.0
GraphIE 2022 [18]	75.7	59.4
Portuguese ACE-2005		
BERT-ACE05PT [6] (ours)	64.4	46.7

4 ERROR ANALYSIS AND LIMITATIONS

Firstly, it is important to acknowledge that automatic translation is prone to translation errors which can lead to inaccuracies and a loss of the intended meaning in the translated texts.

Regarding annotation alignments, some errors were found. For instance, in non-null subject languages like English, it is obligatory to include the subject in sentence construction, while in Portuguese, the subject can be implicit [3]. Consider the next sentence as an example:

"We have been under heavy fire in the Middle East"

In this example, the word "fire" is marked as an event trigger of type Conflict:Attack, and "We" is an event argument with the role Target. However, in the Portuguese translation, "We have" was translated to "*Temos*" (the conjugated verb "have" in the first person plural). As a result, the subject became implicit, omitting the argument "We". During the validation process on the test split of the ACE-2005-PT corpus, we detected 41 errors in the alignment of arguments related to this phenomenon. These errors account for 11.71% of the total argument alignment errors.

Another identified issue consists of our pipeline's inability to accurately identify the sentences' syntax structure. According to the ACE-2005 annotation guidelines [1], when annotating a noun phrase (NP) with a determiner, the latter should be included in the argument. For instance, in the previous example, the NP "*the Middle East*", which is an event argument with the role Place, comprises the determiner "*the*". In this case, the NP is part of a prepositional phrase (PP) headed by the preposition "*in*", but the preposition is not part of the argument. In the context of the Portuguese language, in examples like this one, the preposition is contracted with the determiner, that is, in the PP "*no Médio Oriente*", we observe the contraction of "*em*" ("*in*") + "*o*" ("*the*"). The event argument in Portuguese is "*o Médio Oriente*" and not "*no Médio Oriente*". While manually annotating the test set of the ACE-2005 Portuguese version, our expert followed ACE-2005 guidelines, which meant separating the preposition from the determiner. However, this requirement poses a challenge for automatic alignment.

We detected 242 instances where our pipeline failed to correctly align the determiner, accounting for 69.14% of the errors in argument alignment. This is one of the main reasons for the discrepancy between the relaxed (87.77%) and Exact (70.55%) scores. By using

the exact match, these cases are considered completely misaligned. On the other hand, by using the relaxed match, despite being penalised for omitting the determiner, the pipeline still receives credit for identifying the remaining correctly aligned tokens.

We also conducted an error analysis on the pipeline text alignment components. For instance, character-level similarity matching disregards the order of characters and lacks semantic understanding. In the previous example, the argument annotation "We" was translated to "*Nós*" in isolation. By relying on character similarity matching, "*Nós*" was mistakenly aligned with the word "*no*" (the) due to their shared characters, resulting in an incorrect alignment. This issue becomes more prevalent when dealing with shorter strings, thus, this method was not used to align triggers which are usually composed of one single word. On the other hand, the lemmatization method also introduces ambiguity, as multiple words with different meanings can map to the same lemma. This ambiguity may result in incorrect alignments, where annotations are aligned with incorrect spans in the translated text due to multiple possible lemma matches.

In Algorithm 3, we determine the translated annotation offsets using the maximum and minimum indexes of the previously computed alignments. However, in case of alignment errors, the calculated minimum or maximum values can become disproportionately large or small, resulting in absurdly large annotation spans. To address this issue, we implemented a safeguard by comparing the lengths of the aligned annotation and the source annotation. If a significant discrepancy is detected, the alignment candidate is discarded.

5 CONCLUSION

In this work, we developed a translation and alignment pipeline specifically designed to automatically translate the ACE-2005 corpus into both European and Brazilian Portuguese. To assess its effectiveness, we compared the pipeline results against manually aligned annotations performed by a linguistic expert, achieving a relaxed score of 87.77% and an exact match of 70.55%. While our primary focus was the Portuguese language, our pipeline can be easily adapted to accommodate other languages and corpora. This work expanded the usability of the ACE-2005 corpus beyond its original three languages by creating a Portuguese version, which has been accepted for publication by the LDC.

In the future, our pipeline could translate ACE-2005 into other languages and extend to other corpora such as the Entities, Relations and Events corpus (ERE) [5], enhancing NLP tasks in various languages and tasks.

6 ACKNOWLEDGMENTS

This work is financed by National Funds through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia, within project LA/P/0063/2020. DOI 10.54499/LA/P/0063/2020 | <https://doi.org/10.54499/LA/P/0063/2020>. The author also would like to acknowledge the project StorySense, with reference 2022.093 12.PTDC (DOI 10.54499/2022.09312.PTDC).

REFERENCES

- [1] 2005. English Annotation Guidelines for Events. *Linguistic Data Consortium* (2005). <https://www.ldc.upenn.edu/sites/www.ldc.upenn.edu/files/english->

- events-guidelines-v5.4.3.pdf
- [2] Ali Balali, Masoud Asadpour, Ricardo Campos, and Adam Jatowt. 2020. Joint event extraction along shortest dependency paths using graph convolutional networks. *Knowledge-Based Systems* 210 (2020), 106492. <https://doi.org/10.1016/j.knsys.2020.106492>
- [3] Pilar Barbosa, Maria Eugênia L Duarte, and Mary Aizawa Kato. 2005. Null subjects in European and Brazilian Portuguese. *Journal of Portuguese Linguistics* 4, 2 (2005).
- [4] Casimiro Pio Carrino, Marta R Costa-Jussà, and José A R Fonollosa. 2020. Automatic Spanish Translation of the SQuAD Dataset for Multilingual Question Answering. (2020), 11–16.
- [5] Song Chen, Ann Bies, Kira Griffitt, Joe Ellis, and Stephanie Strassel. 2023. DEFT English Light and Rich ERE Annotation. <https://doi.org/11272.1/AB2/7KH7V4>
- [6] Luís Filipe Cunha, Ricardo Campos, and Alípio Jorge. 2023. Event Extraction for Portuguese: A QA-driven Approach using ACE-2005. In *Springer's LNAI – Lecture Notes in Artificial Intelligence*. Springer.
- [7] Jurafsky Daniel and James H Martin. 2023. Relation and Event Extraction. In *Speech and Language Processing*. Chapter 21, 429–445.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR* abs/1810.04805 (2018). arXiv:1810.04805
- [9] George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. The Automatic Content Extraction (ACE) Program – Tasks, Data, and Evaluation. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*. European Language Resources Association (ELRA), Lisbon, Portugal.
- [10] Zi Yi Dou and Graham Neubig. 2021. Word Alignment by Fine-tuning Embeddings on Parallel Corpora. *EACL 2021 - 16th Conference of the European Chapter of the Association for Computational Linguistics, Proceedings of the Conference* (2021), 2112–2128.
- [11] Kinya Du and Claire Cardie. 2020. Event Extraction by Answering (Almost) Natural Questions. *EMNLP 2020 - 2020 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference* (2020), 671–683.
- [12] Vladimir I Levenshtein. 1965. Binary Codes Capable of Correcting Deletions, Insertions, and Reversals. 163, 4 (1965), 845–848.
- [13] Qi Li, Heng Ji, and Liang Huang. 2013. Joint Event Extraction via Structured Prediction with Global Features. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Bulgaria, 73–82.
- [14] Qian Li, Jianxin Li, Jiawei Sheng, Shiyao Cui, Jia Wu, Yiming Hei, Hao Peng, Shu Guo, Lihong Wang, Amin Beheshti, and Philip S. Yu. 2022. A Survey on Deep Learning Event Extraction: Approaches and Applications. arXiv:2107.02126 [cs.CL]
- [15] Ying Lin, Heng Ji, Fei Huang, and Lingfei Wu. 2020. A Joint Neural Model for Information Extraction with Global Features. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 7999–8009.
- [16] Yaojie Lu, Hongyu Lin, Jin Xu, Xianpei Han, Jialong Tang, Annan Li, Le Sun, Meng Liao, and Shaoyi Chen. 2021. Text2Event: Controllable Sequence-to-Structure Generation for End-to-end Event Extraction. 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)*. 2795–2806.
- [17] Minh Van Nguyen, Viet Dac Lai, and Thien Huu Nguyen. 2021. Cross-Task Instance Representation Interactions and Label Dependencies for Joint Information Extraction with Graph Convolutional Networks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 27–38. <https://doi.org/10.18653/v1/2021.naacl-main.3>
- [18] Minh Van Nguyen, Bonan Min, Franck Dernoncourt, and Thien Nguyen. 2022. Joint Extraction of Entities, Relations, and Events via Modeling Inter-Instance and Inter-Label Dependencies. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. United States, 4363–4374. <https://doi.org/10.18653/v1/2022.naacl-main.324>
- [19] Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. Joint Event Extraction via Recurrent Neural Networks. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. California, 300–309. <https://doi.org/10.18653/v1/N16-1034>
- [20] John W Ratcliff and David E Metzener. 1988. Pattern-matching—the gestalt approach. *Dr Dobbs Journal* 13, 7 (1988), 46.
- [21] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. arXiv:1706.03762 [cs.CL]