

A logical-based corpus for cross-lingual evaluation*

Felipe Salvatore¹, Marcelo Finger^{1†} and R. Hirata Jr^{1‡}

¹Department of Computer Science, Instituto de Matemática e Estatística,
University of São Paulo, Brazil
{felsal, mfinger, hirata}@ime.usp.br

Abstract

At present, different deep learning models are presenting high accuracy on popular inference datasets such as SNLI, MNLI, and SciTail. However, there are different indicators that those datasets can be exploited by using some simple linguistic patterns. This fact poses difficulties to our understanding of the actual capacity of machine learning models to solve the complex task of textual inference. We propose a new set of syntactic tasks focused on contradiction detection that require specific capacities over linguistic logical forms such as: Boolean coordination, quantifiers, definite description, and counting operators. We evaluate two kinds of deep learning models that implicitly exploit language structure: recurrent models and the Transformer network BERT. We show that although BERT is clearly more efficient to generalize over most logical forms, there is space for improvement when dealing with counting operators. Since the syntactic tasks can be implemented in different languages, we show a successful case of cross-lingual transfer learning between English and Portuguese.

1 Introduction

Natural Language Inference (NLI) is a complex problem of Natural Language Understanding which is usually defined as follows: given a pair of textual inputs P and H we need to determine if P entails H , or H contradicts P , or H and P have no logical relationship (they are *neutral*) The Fracas Consortium et al. (1996). P and H , known

as “*premise*” and “*hypothesis*” respectively, can be either simple sentences or full texts.

The task can focus either on the entailment or the contradiction part. The former, which is known as Recognizing Textual Entailment (RTE) Dagan et al. (2013), classifies the pair P, H in “*entailment*” or “*non-entailment*”. The latter, which is known as Contradiction Detection (CD), classifies that pair in terms of “*contradiction*” or “*non-contradiction*”. Independently of the form that we frame the problem, the concept of inference is the critical issue here.

With this formulation, NLI has been treated as a text classification problem suitable to be solved by a variety of machine learning techniques Bowman et al. (2015a); Williams et al. (2017). Inference itself is also a complex problem. As shown in the following sentence pairs:

1. “*A woman plays with my dog*”, “*A person plays with my dog*”
2. “*Jenny and Sally play with my dog*”, “*Jenny plays with my dog*”

Both examples are cases of entailment, with different properties. In (1) the entailment is caused by the hypernym relationship between “*person*” and “*woman*”. Example (2) deals with interpretation of the coordinating conjunction “*and*” as a Boolean connective. As (1) relies on the meaning of the noun phrases we call it “*lexical inference*”. As (2) is invariant under substitution we call it “*structural inference*”. The latter is the focus of this work.

In this paper, we propose a new synthetic CD dataset that enables us to:

1. compare the NLI accuracy of different deep learning models.

*This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001; and Fapesp 2019/07665-4.

[†]Partly supported by Fapesp project 2014/12236-1 and CNPq grant PQ 303609/2018-4.

[‡]Partly supported by FAPESP projects 2015/01587-0, 2015/24485-9 and 2017/25835-9

2. diagnose the structural (logical and syntactic) competence of each model.
3. verify the cross-lingual performance of each method.

The contributions presented in this paper are: i) the presentation of a structure oriented CD dataset; ii) the comparison of traditional neural recurrent models against the Transformer network BERT; iii) a success case of cross-lingual transfer learning for structural NLI between English and Portuguese.

2 Background and Related Work

The size of NLI datasets has been increasing since the initial proposition of the FraCas test suit composed of 346 examples [The Fracas Consortium et al. \(1996\)](#). Some old datasets like RTE-6 [Bentivogli et al. \(2009\)](#) and SICK [Marelli et al. \(2014\)](#), with 16K and 9.8K examples, respectively, are relatively small if compared with the current ones like SNLI [Bowman et al. \(2015a\)](#) and MNLI [Williams et al. \(2017\)](#), with 570K and 433K examples, respectively. This increase was possible with the use of crowdsourcing platforms like the Amazon Mechanical Turk [Bowman et al. \(2015a\)](#); [Williams et al. \(2017\)](#). The annotation performed by a formal semanticist, like in RTE 1-3 [Giampiccolo et al. \(2007\)](#), was replaced with the generation of sentence pairs done by average English speakers. This change in dataset construction has been criticised with the argument that it is hard for an average speaker to produce different and creative examples of entailment and contradiction pairs [Gururangan et al. \(2018\)](#). By looking at the hypothesis alone a simple text classifier can achieve an accuracy significantly better than a random classifier in datasets such as SNLI and MNLI. This was explained by a high correlation of occurrences of negative words (“no”, “nobody”, “never”, “nothing”) in contradiction instances, and high correlation of generic words (such as “animal”, “instrument”, “outdoors”) with entailment instances. Thus, despite of the large size of the corpora the task was easier to perform than expected [Poliak et al. \(2018\)](#).

The new wave of pre-trained models [Howard and Ruder \(2018\)](#); [Devlin et al. \(2018\)](#); [Liu et al. \(2019\)](#) poses both a challenge and an opportunity for the NLI field. The large-scale datasets are close to being solved (the benchmark for SNLI,

MNLI, and SciTail is 91.1%, 85.3%/85.0%, and 94.1%, respectively, as reported in [Liu et al. \(2019\)](#)), giving the impression that NLI will become a trivial problem. The opportunity lies in the fact that, by using pre-trained models, training will no longer need such large datasets. Then we can focus our efforts in creating small, well-thought datasets that reflect the variety of inferential tasks, and so determine the real competence of a model.

Here we present a collection of small datasets designed to measure the competence of detecting contradictions in structural inferences. We have chosen the CD task because it is harder for an average annotator to create examples of contradictions without excessively relying on the same patterns. At the same time, CD has practical importance since it can be used to improve consistency in real case applications, such as chat-bots [Welleck et al. \(2018\)](#).

We choose to focus on structural inference because we have detected that the current datasets are not appropriately addressing this particular feature. In an experiment, we verify the deficiency reported in [Gururangan et al. \(2018\)](#); [Glockner et al. \(2018\)](#). First, we transformed the SNLI and MNLI datasets to a CD task. The transformation is done by converting all instances of entailment and neutral into non-contradiction, and by balancing the classes in both training and test data. Second, we applied a simple Bag-of-Words classifier, destroying any structural information. The accuracy was significantly higher than the random classifier, 63.9% and 61.9% for SNLI and MNLI, respectively. Even the recent dataset focusing on contradiction, Dialog NLI [Welleck et al. \(2018\)](#), presents a similar pattern. The same Bag-of-Words model achieved 76.2% accuracy in this corpus.

Our approach of isolating structural forms by using synthetic data to analyze the logical and syntactical competence of different neural models is similar to [Bowman et al. \(2015b\)](#); [Evans et al. \(2018\)](#); [Tran et al. \(2018\)](#). One main difference between their approach and ours is that we are interested in using a formal language as a tool for performing a cross-lingual analysis.

3 Data Collection

The different datasets that we propose are divided by tasks, such that each task introduces a new linguistic construct. Each task is designed by applying structurally dependent rules to automatically

generate the sentence pairs. We first define the pairs in a formal language and then we use it to generate instances in natural language. In this paper, we have decided to work with English and Portuguese.

There are two main reasons to use a formal language as a basis for the dataset. First, this approach allows us to minimize the influence of common knowledge and lexical knowledge, highlighting structural features. Second, we can obtain a structural symmetry between the English and Portuguese corpora.

Hence, our dataset is a tool to measure inference in two dimensions: one defined by the structural forms, which corresponds to different levels in our hierarchical corpus; and other defined by the instantiation of these forms in multiple natural languages.

3.1 Template Language

The *template language* is a formal language used to generate instances of contradictions and non-contradictions in a natural language. This language is composed of two basic entities: people, $Pe = \{x_1, x_2, \dots, x_n\}$ and places, $Pl = \{p_1, p_2, \dots, p_m\}$. We also define three binary relations: $V(x, y)$, $x > y$, $x \geq y$. It is a simplistic universe with the intended meaning for binary relations such as “ x has visited y ”, “ x is taller than y ” and “ x is as tall as y ”, respectively.

A *realisation* of the template language r is a function mapping Pe and Pl to nouns such that $r(Pe) \cap r(Pl) = \emptyset$; it also maps the relation symbols and logic operators to corresponding forms in some natural language.

Each task is defined by the introduction of a new structural and logical operator. We define the tasks in a hierarchical fashion: if a logical operator appears on a task n , it can appear in any task k (with $k > n$). The main advantage of our approach compared to other datasets is that we can isolate the occurrences of each operator to have a clear notion in what forces the models to fail (or succeed).

For each task, we provide training and test data with 10K and 1K examples, respectively. All data is balanced; and, as usual, the model’s accuracy is evaluated on the test data. To test the model’s generalization capability, we have defined two distinct realization functions r_{train} and r_{test} such that $r_{train}(Pe) \cap r_{test}(Pe) = \emptyset$ and $r_{train}(Pl) \cap r_{test}(Pl) = \emptyset$. For example, in the English ver-

sion $r_{train}(Pe)$ and $r_{train}(Pl)$ are composed of common English masculine names and names of countries, respectively. Similarly, $r_{test}(Pe)$ and $r_{test}(Pl)$ are composed of feminine names and names of cities from the United States. In the Portuguese version we have done a similar construction, using common masculine and feminine names together with names of countries and names of Brazilian cities.

3.2 Data Generation

A logical rule can be seen as a mapping that transforms a premise P into a conclusion C . To obtain examples of contradiction we start with a premise P and define H as the negation of C . The examples of non-contradiction are different negations that do not necessarily violate P . We repeat this process for each task. What defines the difference from one task to another is the introduction of logical and linguist operators, and subsequently, new rules. We have used more than one template pair to define each task; however, for the sake of brevity, in the description below we will give only a brief overview of each task.

The full dataset in both languages, together with the code to generate it and the detailed list of all templates, can be found online [Salvatore \(2019\)](#).

Task 1: Simple Negation We introduce the negation operator \neg , “not”. The premise P is a collection of facts about some agents visiting different places. Example, $P := \{V(x_1, p_1), V(x_2, p_2)\}$ (“*Charles has visited Chile, Joe has visited Japan*”). The hypothesis H can be either a negation of one fact that appears in P , $\neg V(x_2, p_2)$ (“*Joe didn’t visit Japan*”); or a new fact not related to P , $\neg V(x, p)$ (“*Lana didn’t visit France*”). The number of facts that appear in P vary from two to twelve.

Task 2: Boolean Coordination In this task, we add the Boolean conjunction \wedge , the coordinating conjunction “and”. Example, $P := \{V(x_1, p) \wedge V(x_2, p) \wedge V(x_3, p)\}$ (“*Felix, Ronnie, and Tyler have visited Bolivia*”). The new information H can state that one of the mentioned agents did not travel to a mentioned place, $\neg V(x_3, p)$ (“*Tyler didn’t visit Bolivia*”). Or it can represent a new fact, $\neg V(x, p)$ (“*Bruce didn’t visit Bolivia*”).

Task 3: Quantification By adding the quantifiers \forall and \exists , “for every” and “some”, respectively, we can construct example of inferences that explicitly exploit the difference between the

two basic entities, people and places. Example, P states a general fact about all people, $P := \{\forall x \forall p V(x, p)\}$ (“Everyone has visited every place”). H can be the negation of one particular instance of P , $\neg V(x, p)$ (“Timothy didn’t visit El Salvador”). Or a fact that does not violate P , $\neg V(x, x_1)$ (“Timothy didn’t visit Anthony”).

Task 4: Definite Description One way to test if a model can capture reference is by using definite description, i.e., by adding the operator ι to perform description and the equality relation $=$. Hence, $x = \iota y Q(y)$ is to be read as “ x is the one that has property Q ”. Here we describe one property of one agent and ask the model to combine the description with a new fact. For example, $P := \{x_1 = \iota y \forall p V(y, p), V(x_1, x_2)\}$ (“Carlos is the person that has visited every place, Carlos has visited John”). Two new hypotheses can be introduced: $\neg V(x_1, p)$ (“Carlos did not visit Germany”) or $\neg V(x_2, p)$ (“John did not visit Germany”). Only the first hypothesis is a contradiction. Although the names “Carlos” and “John” appear on the premise, we expected the model to relate the property “being the one that has visited every place” to “Carlos” and not to “John”.

Task 5: Comparatives In this task we are interested to know if the model can recognise a basic property of a binary relation: transitivity. The premise is composed of a collection of simple facts $P := \{x_1 > x_2, x_2 > x_3\}$. (“Francis is taller than Joe, Joe is taller than Ryan”). Assuming the transitivity of $>$, the hypothesis can be a consequence of P , $x_1 > x_3$ (“Francis is taller than Ryan”), or a fact that violates the transitivity property, $x_3 > x_1$ (“Ryan is taller than Francis”). The size of the P varies from four to ten. Negation is not employed here.

Task 6: Counting In Task 3 we have added only the basic quantifiers \forall and \exists , but there is a broader family of operators called *generalised quantifiers*. In this task we introduce the counting quantifier $\exists_{=n}$ (“exactly n ”). Example, $P := \{\exists_{=3} p V(x_1, p) \wedge \exists_{=2} x V(x_1, x)\}$ (“Philip has visited only three places and only two people”). H can be an information consistent with P , $V(x_1, x_2)$ (“Philip has visited John”), or something that contradicts P , $V(x_1, x_2) \wedge V(x_1, x_3) \wedge V(x_1, x_4)$ (“Philip has visited John, Carla, and Bruce”). We have added counting quantifiers corresponding to numbers from one to thirty.

Task 7: Mixed In order to guarantee variability,

we created a dataset composed of different samples of the previous tasks.

Basic statistics for the English and Portuguese realisations of all tasks can be found in Table 1.

Task	Vocab size	Vocab intersection	Mean input length	Max input length
1 (Eng)	3561	77	230.6	459
2 (Eng)	4117	128	151.4	343
3 (Eng)	3117	70	101.5	329
4 (Eng)	1878	62	100.81	134
5 (Eng)	1311	25	208.8	377
6 (Eng)	3900	150	168.4	468
7 (Eng)	3775	162	160.6	466
1 (Pt)	7762	254	209.4	445
2 (Pt)	9990	393	148.5	388
3 (Pt)	5930	212	102.7	395
4 (Pt)	5540	135	91.8	140
5 (Pt)	5970	114	235.2	462
6 (Pt)	9535	386	87.8	531
7 (Pt)	8880	391	159.9	487

Table 1: Task description. Column 1 presents two realizations of the described tasks - one in English (Eng) and the other in Portuguese (Pt). Column 2 presents the vocabulary size for the task. Column 3 presents the number of words that occurs both in the training and test data. Column 4 presents the average length in words of the input text (the concatenation of P and H). Column 5 presents the maximum length of the input text.

Since we are using a large number of facts in P , the input text is longer than the ones presented in average NLI datasets.

4 Models and Evaluation

To evaluate the accuracy of each CD task we employed three kinds of models:

Baseline The baseline model (Base) is a Random Forest classifier that models the input text, the concatenation of P and H , using the Bag-of-Words representation. Since we have constructed the dataset centered on the notion of structure-based contradictions, we believe that it should perform slightly better than random. At the same time, by using such baseline, we can certify if the proposed tasks are indeed requiring structural knowledge.

Recurrent Models The dominant family of neural models in Natural Language Processing specialised in modelling sequential data is the

one composed by the *Recurrent Neural Networks* (RNNs) and its variations, *Long Short-Term Memory* (LSTM), and *Gated Recurrent Unit* (GRU) Goldberg (2015). We consider both the standard and the bidirectional variants of this family of models. As input for these models, we use the concatenation of P and H as a single sentence.

Traditional multilayer recurrent models are not the best choice to improve the benchmark on NLI Glockner et al. (2018). However, in recent works, it has been reported that recurrent models achieve a better performance than Transformer-based models to capture structural patterns for logical inference Evans et al. (2018); Tran et al. (2018). We want to investigate if the same result can be achieved using our tasks as the base of comparison.

Transformer-based Models A recent non-recurrent family of neural models known as *Transformer networks* was introduced in Vaswani et al. (2017). Different from the recurrent models that recursively summarizes all previous input into a single representation, the Transformer network employs a self-attention mechanism to directly attend to all previous inputs (more details of this architecture can be found in Vaswani et al. (2017)). Although, by performing regular training using this architecture alone we do not see surprising results in inference prediction Evans et al. (2018); Tran et al. (2018), when we pre-trained a Transformer network in the language modeling task and fine-tuned afterwards on an inference task we see a significant improvement Devlin et al. (2018).

Among the different Transformer-based models we will focus our analysis on the multilayer bidirectional architecture known as *Bidirectional Encoder Representation from Transformers* (BERT) Devlin et al. (2018). This bidirectional model, pre-trained as a masked language model and as a next sentence predictor, has two versions: BERT_{BASE} and BERT_{LARGE}. The difference lies in the size of each architecture, the number of layers and self-attention heads. Since BERT_{LARGE} is unstable on small datasets Devlin et al. (2018) we have used only BERT_{BASE}.

The strategy to perform NLI classification using BERT is the same the one presented in Devlin et al. (2018): together with the pair P, H we add new special tokens [CLS] (classification token) and [SEP] (sentence separator). Hence, the textual input is the result of the concatenation:

[CLS] P [SEP] H [SEP]. After we obtain the vector representation of the [CLS] token, we pass it through a classification layer to obtain the prediction class (contradiction / non-contradiction). We fine-tune the model for the CD task in a standard way, the original weights are co-trained with the weights from the new layer.

By comparing BERT with other models we are not only comparing different architectures but different techniques of training. The baseline model uses no additional information. The recurrent models use only a soft version of transfer learning with fine-tuning of pre-trained embeddings (the fine-tuning of one layer only). On the other side, BERT is pre-trained on a large corpus as a language model. It is expected that this pre-training helps the model to capture some general properties of language Howard and Ruder (2018). Since the tasks that we proposed are basic and cover very specific aspects of reasoning, we can use it to evaluate which properties are being learned in the pre-training phase.

The simplicity of the tasks motivated us to use transfer-learning differently: instead of simply using the multilingual version of BERT¹ and fine-tune it on the Portuguese version of the tasks, we have decided to check the possibility of transferring structural knowledge from high-resource languages (English / Chinese) to Portuguese.

This can be done because for each pre-trained model there is a tokenizer that transforms the Portuguese input into a collection of tokens that the model can process. Thus, we have decided to use the regular version of BERT trained on an English corpus (BERT_{eng}), the already mentioned Multilingual BERT (BERT_{mult}), and the version of the BERT model trained on a Chinese corpus (BERT_{chi}).

We hypothesize that *most structural patterns learned by the model in English can be transferred to Portuguese*. By the same reasoning, we believe that BERT_{chi} should perform poorly. Not only the tokenizer associated to BERT_{chi} will add noise to the input text, but also Portuguese and Chinese are grammatically different; for example, the latter is overwhelmingly right-branching while the former is more mixed Levy and Manning (2003).

¹Multilingual BERT is a model trained on the concatenation of the entire Wikipedia from 100 languages, Portuguese included. <https://github.com/google-research/bert/blob/master/multilingual.md>

4.1 Experimental settings

Given the above considerations, four research questions arose:

- (i) *How the different models perform on the proposed tasks?*
- (ii) *How much each model rely on the occurrence of non-logical words?*
- (iii) *Can cross-lingual transfer learning be successfully used for the Portuguese realization of those tasks?*
- (iv) *Is the dataset biased? Are the models learning some unexpected text pattern?*

To answer those questions, we evaluated the models performance in four different ways:

- (i) Each model was trained on different proportions of the dataset. In this case, $r_{train}(Pe) \cap r_{test}(Pe) = \emptyset$ and $r_{train}(Pl) \cap r_{test}(Pl) = \emptyset$.
- (ii) We have trained the models on a version of the dataset where we allow full intersection of the train and test vocabulary, i.e., $r_{train}(Pe) = r_{test}(Pe)$ and $r_{train}(Pl) = r_{test}(Pl)$.
- (iii) For the Portuguese corpus, we have fine-tuned the three pre-trained models mentioned previously: BERT_{eng}, BERT_{mult}, and BERT_{chi}.
- (iv) We have trained the best model from (i) on the following modified versions of the dataset:
 - (a) *Noise label* - each pair P, H is unchanged but we randomly labeled the pair as contradiction or non-contradiction.
 - (b) *Premise only* - we keep the labels the same and omit the hypothesis H .
 - (c) *Hypothesis only* - the premise P is removed, but the labels remain intact.

4.2 Implementation

All deep learning architectures were implemented using the Pytorch library Paszke et al. (2017). To make use of the pre-trained version of BERT we have based our implementation on the

public repository <https://github.com/huggingface/pytorch-pretrained-BERT>.

The different recurrent architectures were optimized with Adam Kingma and Ba (2014). We have used pre-trained word embedding from Glove Pennington et al. (2014) and Fasttext Joulin et al. (2016), but we also used random initialized embeddings. We random searched across embedding dimensions in $[10, 500]$, hidden layer size of the recurrent model in $[10, 500]$, number of recurrent layer in $[1, 6]$, learning rate in $[0, 1]$, dropout in $[0, 1]$ and batch sizes in $[32, 128]$.

The hyperparameter search for BERT follows the one presented in Devlin et al. (2018) that uses Adam with learning rate warmup and linear decay.

We randomly searched the learning rate in $[2 \cdot 10^{-5}, 5 \cdot 10^{-5}]$, batch sizes in $[16, 32]$ and number of epochs in $[3, 4]$.

All the code for the experiments is public available Salvatore (2019).

4.3 Results

How the different models perform on the proposed tasks?

In most of the tasks, BERT_{eng} presents a clear advantage when compared to all other models. Tasks 3 and 6 are the only ones where the difference in accuracy between BERT_{eng} and the recurrent models is small, as can be seen in Table 2. Even when we look at BERT_{eng}'s results on the Portuguese corpus, which are slightly worse when compared to the English one, we still see a similar pattern.

Figure 1 shows that BERT_{eng} is the only model improved by training on more data. All other models remain close to random independently of the amount of training data.

Accuracy improvement over training size indicates the difference in difficulty of each task. On the one hand, Tasks 1, 2 and 4 are practically solved by BERT using only 4K examples of training (99.5%, 99.7%, 97.6% accuracy, respectively). On the other hand, the results for Tasks 3 and 6 remain below average, as seen in Figure 2.

How much each model rely on the occurrence of non-logical words?

With the full intersection of the vocabulary, experiment (ii), we have observed that the average accuracy improvement differs from model to model: Baseline, GRU, BERT_{eng}, LSTM and RNN present an average improvement of 17.6%,

Task	Base	RNN	GRU	LSTM	BERT
1 (Eng)	52.1	50.1	50.6	50.4	99.8
2 (Eng)	50.7	50.2	50.2	50.8	100
3 (Eng)	63.5	50.3	66.1	63.5	90.5
4 (Eng)	51.0	51.7	52.7	51.6	100
5 (Eng)	50.6	50.1	50.2	50.2	100
6 (Eng)	55.5	84.4	82.7	75.1	87.5
7 (Eng)	54.1	50.9	53.7	50.0	94.6
Avg.	53.9	55.4	58.0	56.2	96.1
1 (Pt)	53.9	50.1	50.2	50.0	99.9
2 (Pt)	49.8	50.0	50.0	50.0	99.9
3 (Pt)	61.7	50.0	70.6	50.1	78.7
4 (Pt)	50.9	50.0	50.4	50.0	100
5 (Pt)	49.9	50.1	50.8	50.0	99.8
6 (Pt)	58.9	66.4	79.7	67.2	79.1
7 (Pt)	55.4	51.1	51.6	51.1	82.7
Avg.	54.4	52.6	57.6	52.6	91.4

Table 2: Results of the experiment (i), accuracy percentage on test data for the English and Portuguese corpora

9.6%, 5.3%, 4.25%, 1.3%, respectively. This may indicate that the recurrent models are relying more on noun phrases than BERT. However, since the difference is not significant, more investigation is required.

Can cross-lingual transfer learning be successfully used for the Portuguese realization of those tasks?

As expected, when we fine-tuned BERT_{multi} to the Portuguese version of the dataset we have observed an overall improvement. Most notably, in Tasks 6 and 7 we have achieved a new accuracy of 87.4% and 92.3% respectively. Surprisingly, BERT_{chi} is able to solve some simple tasks, namely Tasks 1, 2 and 4. But when trained on the mixed version of the dataset, Task 7, this pre-trained model had repeatedly present a random performance.

One of the most important features observed by evaluating the different pre-training models is that although BERT_{eng} and BERT_{mult} show a similar result on the Portuguese corpus, BERT_{eng} needs more data to improve its performance, as seen in Figure 3.

Is the dataset biased? Are the models learning some unexpected text pattern?

By taking BERT_{eng} as the best classifier, we repeated the training using all the listed data modifi-

cation techniques. The results, as shown in Figure 4, indicate that BERT_{eng} is not memorizing random textual patterns, neither excessively relying on information that appears only in the premise P or the hypothesis H . When we applied it on these versions of the data, BERT_{eng} behaves as a random classifier.

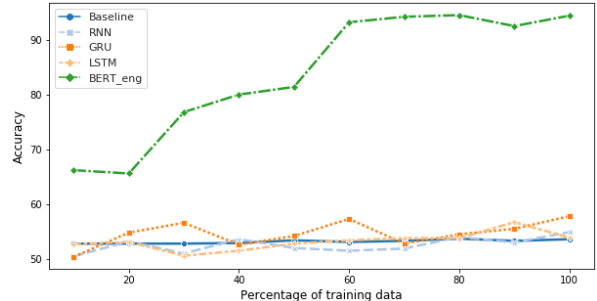


Figure 1: Results of the experiment (i), accuracy for each model on different data proportions (English corpus)

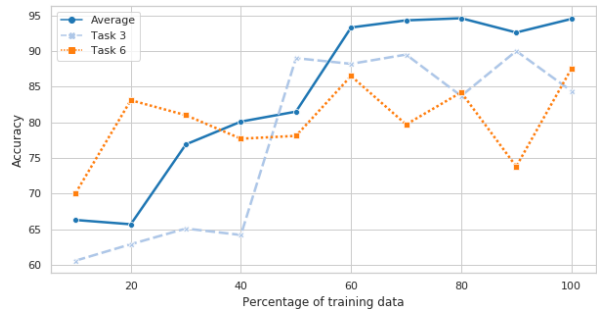


Figure 2: Results of the experiment (i), BERT_{eng}'s accuracy on the different different tasks (English corpus)

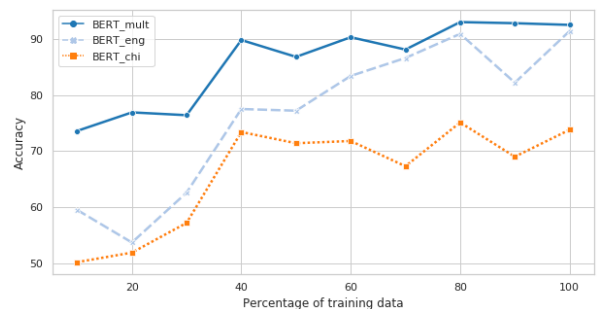


Figure 3: Results of the experiment (iii), different pre-trained BERT versions tested on Portuguese corpus

5 Discussion

The results presented above are similar to the ones reported in Goldberg (2019): *Transformer-based models like BERT can successfully capture syntactic regularities and logical patterns.*

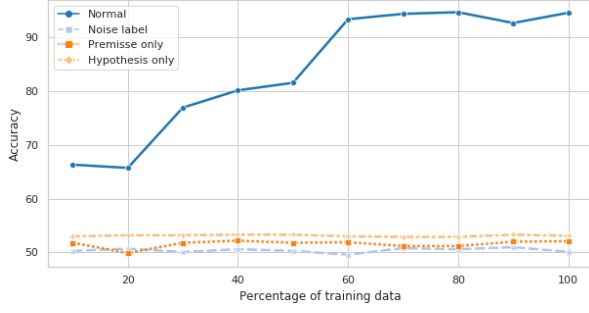


Figure 4: Results of the experiment (iv), $BERT_{eng}$'s accuracy on the different versions of the data (English corpus)

These findings do not contradict the results reported on Evans et al. (2018); Tran et al. (2018), because in both papers, the Transformer models are trained from scratch, while here we have used models that were pre-trained on large datasets with the language model objective.

The results presented both in Table 2 and Figure 3 seem to confirm our initial hypothesis on the effectiveness of transfer learning in a cross-lingual fashion. What has surprised us was the excellent results regarding Tasks 1, 2 and 4 when transferring structural knowledge from Chinese to Portuguese. We offer the following explanation for these results. Take the contradiction pair defined in the template language:

$$P := \{x_1 = \iota y \forall x_2 V(y, x_2), V(x_1, x_3)\}$$

(" x_1 is the person that has visited everybody, x_1 has visited x_3 ")

$$H := \neg V(x_1, x_4) \text{ ("}x_1 \text{ didn't visit } x_4\text{")}$$

If we take one possible Portuguese realization of the pair above and apply the different tokenizers we have the following strings:

1. Original sentence: "[CLS] **gabrielle** é a **pe** soa que visitou **todo mundo** **gabrielle** visitou **luís** [SEP] **gabrielle não** visitou **iane** sis [SEP]".
2. Multilingual tokenizer: "[CLS] **gabrielle** a **pe**soa que visito **##u** **todo mundo** **gabrielle** visito **##u** **lu** **##s** [SEP] **gabrielle** **no** visito **##u** **ian** **##esis** [SEP]".
3. English tokenizer: "[CLS] **gabrielle** a **pe** **##sso** **##a** que visit **##ou** **tod** **##o** **mundo** **gabrielle** visit **##ou** **lu** **##s** [SEP] **gabrielle** **no** visit **##ou** **ian** **##esis** [SEP]".

4. Chinese tokenizer: "[CLS] **ga** **##b** **##rie** **##lle** a **pe** **##ss** **##oa** **q** **##ue** **vi** **##sit** **##ou** **to** **##do** **mu** **##nd** **##o** **ga** **##b** **##rie** **##lle** vi **##sit** **##ou** **lu** **##s** [SEP] **ga** **##b** **##rie** **##lle** **no** vi **##sit** **##ou** **ian** **##es** **##is** [SEP]".

Although the Portuguese words are destroyed by the tokenizers, the model is still able to learn in the fine-tuning phase the *simple* structural pattern between the tokens highlighted above. This may explain why the counting task (Task 4) presents the highest difficulty for BERT. There is some structural grounding for finding contradictions in counting expressions, but to detect contradiction in all cases one must fully grasp the *meaning* of the multiple counting operators.

6 Conclusion

With the possibility of using pre-trained models we can successfully craft small datasets ($\sim 10K$ sentences) to perform fine grained analysis on machine learning models. In this paper, we have presented a new dataset that is able to isolate a few competence issues regarding structural inference. It also allows us to bring to the surface some interesting comparisons between recurrent neural networks and pre-trained Transform-based models. As our results show, *compared to the recurrent models, BERT presents a considerable advantage in learning structural inference. The same result appears even when fine-tuned one version of the model that was not pre-trained on the target language.*

By the stratified nature of our dataset, we can pinpoint BERT's inference difficulties: *there is space for improving the model's counting understanding.* Hence, we can either craft a more realistic NLI dataset centered on the notion of counting or modify BERT's training to achieve better results in the counting task.

The results on cross-lingual transfer learning are stimulating. One possible area for future research is to check if the same results can be attainable using simple structural inferences that occur within complex sentences. This can be done by carefully selecting sentence pairs in a cross-lingual NLI corpus like Conneau et al. (2018). We plan to explore these paths in the future.

References

- Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. 2009. The sixth pascal recognizing textual entailment challenge. In *Text Analysis Conference*.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015a. A large annotated corpus for learning natural language inference. In *Empirical Methods in Natural Language Processing, 2015*.
- Samuel R. Bowman, Christopher D. Manning, and Christopher Potts. 2015b. Tree-structured composition in neural networks without tree-structured architectures. *CoRR*, abs/1506.04834.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. XNLI: Evaluating cross-lingual sentence representations. Association for Computational Linguistics.
- Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzotto. 2013. *Recognizing Textual Entailment: Models and Applications*. Synthesis Lectures on Human Language Technologies. Morgan Claypool Publishers.
- 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.
- Richard Evans, David Saxton, David Amos, Pushmeet Kohli, and Edward Grefenstette. 2018. Can neural networks understand logical entailment? *CoRR*, abs/1802.08535.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. The third PASCAL recognizing textual entailment challenge. In *Proceedings of the Workshop on Textual Entailment and Paraphrasing, Association for Computational Linguistics, 2007*, pages 1–9.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. Breaking NLI systems with sentences that require simple lexical inferences. *CoRR*, abs/1805.02266.
- Yoav Goldberg. 2015. A primer on neural network models for natural language processing. *CoRR*, abs/1510.00726.
- Yoav Goldberg. 2019. Assessing bert’s syntactic abilities. *CoRR*, abs/1901.05287.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R. Bowman, and Noah A. Smith. 2018. Annotation artifacts in natural language inference data. *CoRR*, abs/1803.02324.
- Jeremy Howard and Sebastian Ruder. 2018. Fine-tuned language models for text classification. *CoRR*, abs/1801.06146.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. Bag of tricks for efficient text classification. *CoRR*, abs/1607.01759.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Roger Levy and Christopher Manning. 2003. Is it harder to parse chinese, or the chinese treebank? In *Association for Computational Linguistics, 2003*.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. Multi-task deep neural networks for natural language understanding. *CoRR*, abs/1901.11504.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *LREC*.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing, 2014*.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. Hypothesis only baselines in natural language inference. *CoRR*, abs/1805.01042.
- Felipe Salvatore. 2019. Cross-Lingual Contradiction Detection. <https://github.com/felipessalvatore/CLCD>.
- The Fracas Consortium, Robin Cooper, Dick Crouch, Jan Van Eijck, Chris Fox, Josef Van Genabith, Jan Jaspars, Hans Kamp, David Milward, Manfred Pinkal, Massimo Poesio, Steve Pulman, Ted Briscoe, Holger Maier, and Karsten Konrad. 1996. Using the framework.
- Ke M. Tran, Arianna Bisazza, and Christof Monz. 2018. The importance of being recurrent for modeling hierarchical structure. *CoRR*, abs/1803.03585.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *CoRR*, abs/1706.03762.
- Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2018. Dialogue natural language inference. *CoRR*, abs/1811.00671.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *CoRR*, abs/1704.05426.