

Risks of misinterpretation in the evaluation of distant supervision for relation extraction

Garcia-Mendoza, Juan Luis; Villasenor-Pineda, Luis; Orihuela-Espina, Felipe

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Garcia-Mendoza, JL, Villasenor-Pineda, L & Orihuela-Espina, F 2022, 'Risks of misinterpretation in the evaluation of distant supervision for relation extraction', *Procesamiento de Lenguaje Natural*, vol. 68, pp. 71-83. <<http://journal.sepln.org/sepln/ojs/ojs/index.php/pln/article/view/6407>>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Risks of misinterpretation in the evaluation of Distant Supervision Relation Extraction.

Riesgos de interpretación errónea en la evaluación de la Extracción de Relaciones con Supervisión Distante

Abstract: Distant Supervision is frequently used for addressing Relation Extraction. The evaluation of Distant Supervision in Relation Extraction has been attempted through Precision-Recall curves and/or calculation of Precision at N elements. However, such evaluation is challenging because the labeling of the instances is the result of an automatic process. Consequently, the labels are not necessarily correct, affecting not only the learning process but also the interpretation of the evaluation results. This research aims to show that, should the correct labels be used during the evaluation, the algorithmic performance measured with the mentioned evaluation strategies varies significantly, thus questioning the current interpretation. To this end, we manually labeled a subset of a well-known data set and evaluated the performance of 6 traditional distant supervision approaches. We demonstrate quantitative differences in the evaluation scores when considering manually versus automatically labeled subsets. Consequently, the order of performance among distant supervision algorithms is different.

Keywords: relation extraction, distant supervision evaluation, Precision-Recall curves

Resumen: La Supervisión Distante se utiliza con frecuencia para abordar la extracción de relaciones. La evaluación de la Supervisión Distante en la Extracción de Relaciones se ha realizado mediante curvas de Precisión-Recuerdo y/o el cálculo de la Precisión en N elementos. Sin embargo, dicha evaluación es un desafío porque el etiquetado de las instancias es el resultado de un proceso automático. En consecuencia, las etiquetas no son necesariamente correctas, afectando no solo el proceso de aprendizaje sino también la interpretación de los resultados de la evaluación. El objetivo de esta investigación es mostrar que, si se utilizan las etiquetas correctas durante la evaluación, el desempeño algorítmico medido con las estrategias de evaluación mencionadas varía de manera significativa, cuestionando así la interpretación actual de los resultados que se obtienen. Con este fin, etiquetamos manualmente un subconjunto de un conjunto de datos y evaluamos el desempeño de 6 enfoques tradicionales de Supervisión Distante. Demostramos diferencias cuantitativas en los puntajes de evaluación al considerar subconjuntos etiquetados manualmente versus automáticamente. En consecuencia, el orden de desempeño entre los algoritmos de supervisión distante es diferente.

Palabras clave: extracción de relaciones, evaluación de la supervisión distante, curvas de Precisión-Recuerdo

1 Introduction

Relation Extraction (RE) is concerned with detecting and classifying predefined relations between entities identified in text (Piskorski and Yangarber, 2013). The traditional RE approach uses a supervised method to create the classifier(s) necessary to identify relations between named entities pairs (Hearst, 1992; Agichtein and Gravano, 2000; Bunescu and Mooney, 2005). However, this process is slow and expensive; hence an alternative is the use of Distant Supervision (DS).

DS consists of automatically labeling the relations between each named entities pair in a text using some pre-existing Knowledge Base (KB) (Mintz et al., 2009). It is assumed that in the analyzed text, when two named entities appear observed in the KB, it is assumed that they are associated through the relation existing in the KB. For the automatic annotation of the data set with labeled relations, Mintz et al. (Mintz et al., 2009) assumed that given two entities that participate in a relation, *all* sentences in the data set that includes these two entities express that relation. However, it is not uncommon that a pair of entities in a sentence does not necessarily express a relation or may express several relations (see Fig 1). Hence, Mintz strong assumption often introduces false positives (noise in the labels) in the training and test sets. Later, Riedel et al. (Riedel, Yao, and McCallum, 2010) relaxed this assumption, instead of assuming that “if two entities participate in a relation, *at least one* sentence that mentions these two entities might express that relation”. This relaxation alleviates the problem of false positives in the *automatically* generated labels, but it does not fully fix it.

Unfortunately, the evaluation of these methods is complicated because there is no set correctly labeled to check their performance. This is why alternative evaluation methods have been proposed, such as the Precision-Recall (PR) curves or Precision at N (P@N) elements. However, these measures are calculated using data labeled with the same automatic process; that is, the labels are not necessarily correct, impairing the calculation of the evaluation results.

This paper aims to analyze the use of these evaluation measures showing that when the methods are evaluated using a correctly labeled set, the performance of the algorithms

for DS reported so far varies substantially, thus questioning the current interpretation of the evaluation methods. We assessed the performance of 6 DS algorithms with PR curves and P@N analysis, with a correctly labeled set and with automatically generated labels, and compared the outcomes.

Our contributions can be summarized as follows:

- PR curves and Precision@N performance measures are critically revisited under competing scenarios of manual and automatic labelling.
- The test partition of the *New York Times* (NYT2010)¹ data set proposed by (Riedel, Yao, and McCallum, 2010) was crowd-labelled using MTurk². We argued that this afford better guarantees over the performance assessment in this task.
- We show that under current practice, performance measures for DS in RE may be misinterpreted when evaluation is carried out over automatic potentially noisy labelling.

2 Related Work

The state-of-the-art in DS includes several solutions using different Deep Learning architectures. One of the first networks was the *Piecewise Convolutional Neural Networks* (PCNN) proposed by Zeng et al. (Zeng et al., 2015) based on *Convolutional Neural Networks* (CNN) (Zeng et al., 2014). This network incorporates bags of sentences to handle noise on the labels. A bag of sentences contains sentences that have the same entities pair. Also, it contains a piecewise max pooling layer “to capture structural information between two entities”. Later, different attention mechanisms were incorporated into these architectures. In (Lin et al., 2016; Ji et al., 2017) an attention mechanism at sentences level (CNN_ATT and PCNN_ATT) in multiple instances was proposed to use the information of all sentences in the bag. Also, in (Ji et al., 2017) information about entities was included. Zhou et al. (Zhou et al., 2018) select from the bag several instances related to the label to predict the relations and use a word-level attention mechanism to dynamically highlight important parts of the sentence. Besides, in (Jat, Khandelwal, and

¹<http://iesl.cs.umass.edu/riedel/ecml/>

²Mechanical Turk, MTurk, is a human annotation service provided by Amazon.

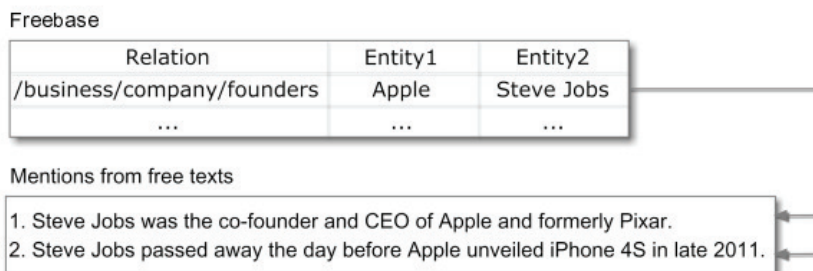


Figure 1: In this example, two sentences with the same pair of entities are automatically labeled with the same relation. Considering the *founders* relation, the first one will be correctly labeled while the second will not (Zeng et al., 2015).

Talukdar, 2018), the *Bidirectional Gated Recurrent Unit* architecture was proposed with an attention mechanism over words to identify which key phrases are used (BGWA). Ye and Ling (Ye and Ling, 2019) used intra-bag and inter-bag attention mechanisms while in (Lin et al., 2016; Ji et al., 2017) it is only performed intra-bag, which ignores when all sentences in the bag are false positives. Moreover, Vashishth et al. (Vashishth et al., 2018) (RESIDE) used knowledge base information such as the entity type and relations alias to predict the correct relation. In addition, *Convolutional Graph Networks* (Deferrard, Bresson, and Vandergheynst, 2016) are used over dependency tree for modeling the syntactic information and capturing long-range dependencies. This information together with the words and positions embeddings is used to encode the entire sentence. On the other hand, Bastos et al. (Bastos et al., 2020) proposed a method using an aggregator that obtains a homogeneous representation with a Graph Neural Network. This representation merges information from sentence, relation and the two entities (considering attributes like entity label, entity alias, entity description and entity type). Many of these methods have been evaluated with the test partition of NYT2010 data set. This partition was automatically labeled under some (imperfect) heuristics and consequently some instances have been associated to an incorrect label. Given the absence of an adequate gold standard, precision, recall, and F1 measurements have not been used for the evaluation of these methods. Mintz et al. (Mintz et al., 2009) used, for the first time, the PR Curves and P@N measures in an attempt to evaluate the DS task. These authors stated that PR curves “gives a rough measure of

precision without requiring expensive human evaluation, making it useful for parameter setting”. But “rough” is not an exact statement; and therefore performance measured with PR curves is dependent on the amount and distribution of noise in the labels. These curves constructed from automatic labels are a simple approximation of the performance of DS methods. Despite this problem, several authors (Surdeanu et al., 2012; Zeng et al., 2015; Lin et al., 2016; Jat, Khandelwal, and Talukdar, 2018; Vashishth et al., 2018; Wu, Fan, and Zhang, 2019; Xu and Barbosa, 2019; Ye and Ling, 2019) continued to use PR curves to evaluate and compare the performance of the proposed DS methods, probably leading to misinterpretations. On the other hand, P@N has been used in DS with 100, 200, 300 and 500 as the value of N. In P@N, the first N elements represent the most reliable answers of the classifier based on the ranking score. Lin et al. (Lin et al., 2016) and Liu et al. (Liu et al., 2017) reported P@100, P@200 and P@300 by randomly extracting one sentence for each pair of entities, two sentences or using them all. This evaluation, like in (Mintz et al., 2009), must be done manually on each execution because of the noise inherent to the *automatic* labels. Unfortunately, many works (Lin et al., 2016; Liu et al., 2017; Wu, Fan, and Zhang, 2019; Vashishth et al., 2018; Ye and Ling, 2019) did not explicitly report whether and how the review was done manually.

Because of the noise that *automatic* labeling introduces, several efforts have been made to build a *gold standard* to evaluate the DS task. First, Mintz et al., (Mintz et al., 2009) used Amazon’s Mechanical Turk service for manual evaluation of P@N. For this, the first 100 instances of each of the top 10 relations

were sent to Mechanical Turk. Hoffmann et al. (Hoffmann et al., 2011) manually labeled 1000 sentences from NYT2010 data set to report the results of their method. These authors stated that “These results provide a good approximation to the true precision but can overestimate the actual recall, since we did not manually check the much larger set of sentences where no approach predicted extractions”. Later, based on these 1000 annotated instances, in (Ren et al., 2017) 395 were used as test partition. However, in these instances there is no more than one sentence per entity pair (Jia et al., 2019). A disadvantage of these data sets is that they do not include the entire NYT2010 test partition. Furthermore, in these papers the measures of the DS task (i.e. PR curves and P@N) were not studied, statistical validations were not carried out, nor was it expressed in which way the selection of the instances was carried out. Finally, with the exception of Hoffmann et al. (Hoffmann et al., 2011), precision, recall and F1 measurements were not reported.

3 Background

3.1 Precision-Recall curves

PR curves are frequently used in binary classification (Davis and Goadrich, 2006) and, within this generic problem, in Information Retrieval (IR) (Manning, Raghavan, and Schütze, 2008). PR curves plot precision versus recall for a varying decision threshold parameter in binary classification (Keilwagen, Grosse, and Grau, 2014). These curves are calculated from the (assumed) true label and a score given by the classifier. This analysis is closely related to the Receiver-Operator Curve (ROC) analysis (Davis and Goadrich, 2006) widely used in statistics. But conveniently for IR purposes, the PR curves can be built without the true negatives TN. To get a scalar score, the area under PR curves (AUC) can be calculated by using the composite trapezoidal method (Davis and Goadrich, 2006).

Let Γ a threshold set defined over classifier scores and Ψ a vector of descending ordered scores given by a classifier. The precision and recall for a threshold $\gamma \in \Gamma$ are calculate by the equations 1 and 2 respectively $\forall \psi \in \Psi \mid \psi > \gamma$.

$$P_\gamma = \frac{TP_\gamma}{TP_\gamma + FP_\gamma} \quad \gamma \in \Gamma \quad (1)$$

$$R_\gamma = \frac{TP_\gamma}{TP_\gamma + FN_\gamma} \quad \gamma \in \Gamma \quad (2)$$

where TP are positive examples correctly labeled as positives, FP are negative examples mislabelled as positives and FN are positive examples incorrectly labeled as negative.

To obtain the set of pairs (R_γ, P_γ) in the PR curve, we iterate over Γ as per Equation 3:

$$PR_Curve(\gamma) = \{(R_\gamma, P_\gamma) : \gamma \in \Gamma\} \quad (3)$$

3.2 Precision at N

The P@N in Equation 4 measures the number of correct elements in a window of N elements (Manning, Raghavan, and Schütze, 2008).

$$P@N = \frac{|TP \cap R_N|}{N} \quad (4)$$

The TP (positive examples correctly labeled as positives) is calculated by manual evaluation. The P@N is frequently used in IR to measure the precision in a subset of retrieved elements R_N , with N the cardinality of the set. According to (Manning, Raghavan, and Schütze, 2008), it has the advantage of not requiring any estimate of the size of the set of relevant elements. While P@N has been used in DS by multiple authors (Zeng et al., 2015; Lin et al., 2016; Ji et al., 2017; He et al., 2018; Wang et al., 2018; Wu, Fan, and Zhang, 2019; Ye and Ling, 2019), but in all these cases, this has been on the automatically labeled data set (with noisy labels).

4 Methodology

4.1 Dataset preparation

In order to establish whether there are risks of misinterpreting the assessment measures, we compared the performance of 6 DS algorithms assessed over *manually-generated* labels and *automatically-generated* labels. We depart from the NYT2010 data set for DS task. This data set includes 53 relations types, including \mathcal{NA} , when there is no relation. Originally, this data set was labeled *automatically*. The *train* partition has 522611 instances (sentence that may or may not contain a relation), 279226 unique entity pairs and 136379 instances with a relation other than \mathcal{NA} . In turn, the test partition has 172448 instances, 96678 unique entity pairs and 6444 instances with a relation other than \mathcal{NA} . From this last partition, in this work,

two test partitions with *manual* labels were built.

In the first test partition, 430 instances were selected for manual revision. The instances selection to be reviewed was made by choosing one instance from each relation at random during 20 iterations. During the manual revision, 88 duplicate instances were found, and 18 that did have unclear relations were detected. Thus, 324 instances were revised *manually* and constitute our first test partition (named *test_1*). Of the 324 instances of the *test_1* partition, 158 changed their *automatic* label after review, i.e. they were considered by a human to hold incorrect labels.

In the second test partition, the complete 6444 instances different from the relation $\mathcal{N}\mathcal{A}$ were selected for manual revision. First, we curated the 6444 instances by removing invalid instances. An invalid instance is considered when the defined entities are not found in the sentence. A total of 6431 were found valid. Then, from the 6431 valid instances we further eliminated 579 instances that contained the same sentence, entity pair, and relation. The rest, 5852, we publish them on the MTurk for review by three reviewers. Finally, we consider an instance as noisy if at least two of the three judges decided that the relations was not expressed. 4801 instances did not vary their *automatic* label but 1051 did (17.9%). This partition we named *test_2*.

4.2 Selection of DS methods for comparison

The following DS methods were compared in their performance:

- PCNN (Zeng et al., 2015) and CNN: The authors used both, PR curves and P@N for evaluation, and labeling was done manually. This was one of the first architectures to be used in DS.
- PCNN_ATT (Lin et al., 2016) and CNN_ATT: The authors incorporated an attention mechanism over instances. They used PR curves to determine the performance of the attention mechanism compared to other methods. Finally, P@N was calculated on automatically generated *automatic* labels.
- BGWA (Jat, Khandelwal, and Talukdar, 2018): It incorporates an attention mechanism over words and entities. In the original work, only the PR curves were used

as the measure of performance.

- RESIDE (Vashishth et al., 2018): It combines syntactic information with entity types and relations aliases. In the original work, it obtained a higher PR curve than previous methods. Like (Lin et al., 2016), P@N was calculated automatically on *automatic labels*.

The main selection criterion for these methods was that they use three different architectures. On the one hand, CNN and PCNN use a convolutional architecture to which an attention mechanism is then incorporated (CNN_ATT and PCNN_ATT). On the other hand, RESIDE uses Graph Convolution Networks and Bidirectional Gated Recurrent Unit and incorporates information on entities and relations. The execution of these methods was done in the same way as defined in Github³ without using the gradient descent optimizer.

They were trained with the NYT2010 train partition proposed by (Riedel, Yao, and McCallum, 2010). The evaluation was carried out for the *test_1* and *test_2* with the *automatic* and *manual* labels (see Figure 2).

4.3 Experimental design

In order to evaluate performance fairly, replications are necessary to ensure that chance does not play a role in our results. The number of replications (sample size) was determined using power analysis. Power analysis refers to the estimation of the probability of correctly rejecting a false null hypothesis when a particular alternative hypothesis is true (Howell, 2012). The analysis depends on four factors: statistical significance, effect size, sample size and the statistical power itself. Fixing any three, yields the fourth for a given hypothesis model. The power analysis was estimated using ANOVA One Way test for a desired significance level of 0.05, statistical power of $\beta = 0.95$ and assuming an effect size of Cohen’s $d = 0.4$. As a result, 42 repetitions per treatment (i.e. algorithm to be compared) was obtained as the required sample size. The number of samples here represents the number of executions for each of the methods, that is, the replications required to detect an effect of the assumed size in the experiment.

From the replications results, the Friedman test was used to determine if there

³<https://github.com/mallabiisc/RESIDE>

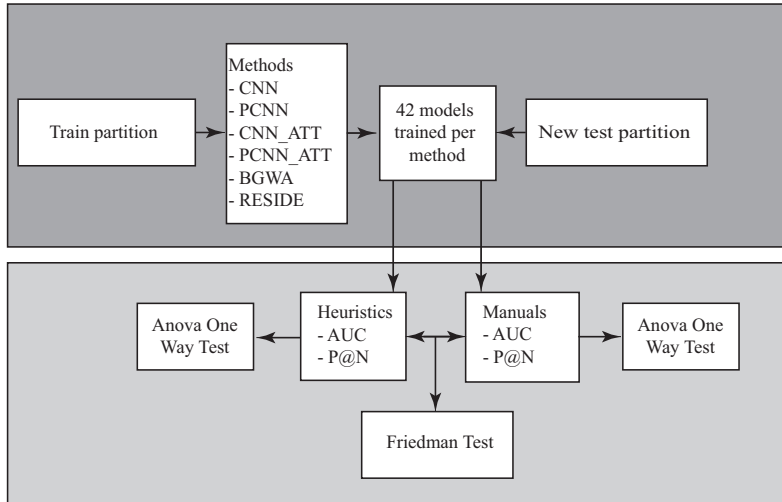


Figure 2: Methodological diagram of this research. The top box illustrates the experiment design. The bottom box summarizes the statistical hypothesis testing.

were differences in the methods ranking using automatic labels with respect to manual labels. The Friedman test is used for one-way repeated measures analysis of variance by ranks (Friedman, 1940). Then, the ANOVA One Way test is applied on *automatic* and *manual* labels to know if there are significant differences between the results achieved by the methods. The ANOVA One Way test is used to test for differences among at least three groups, with the two-group case covered by the simpler *t*-test (Student, 1908; Howell, 2012). Finally, if there were significant differences, pairwise comparisons were made to observe which pair of methods showed differences. The two-by-two comparisons were made with *t*-test and Holm Correction (Holm, 1979). Significance threshold was set at $p < 0.05$.

5 Experiments

5.1 Precision-Recall curves

Performance on *test_1* partition

The Table 1 summarizes the AUC of the tested algorithms PR curves with *automatic* and *manual* labels on *test_1* respectively⁴. All methods increased their AUC with the *manual* labels with regards to their performances using the *automatic* ones, pointing to a systematic overall underestimation. Further, and more critically here, the order of the methods in terms of their performance varied significantly (Friedman: $\chi^2(2) = 373.46$, $p < 2.2e^{-16}$), i.e. they are all underesti-

mated but not in the same extent. This suggests that using PR curves with *automatic* labels might not conferring the direct message one would expect otherwise in the DS evaluation task, and that for this scenario, such bias has to be considered during interpretation. Further, significant differences were found with either *automatics* (ANOVA: $F(5, 246) = 746.9$, $p < 2e^{-16}$) and *manual* labels (ANOVA: $F(5, 246) = 520.8$, $p < 2e^{-16}$).

The Figures 3a and 3b show the PR curves of the methods BGWA, RESIDE, PCNN, PCNN_ATT, CNN and CNN_ATT in one of the executions made with *automatic* and *manual* labels respectively on *test_1*. It can be appreciated that the ordering of the algorithms according to their performance in terms of the area under curve (AUC) varies when using the *manual* labels with respect to the *automatic* ones (previously validated with Friedman test and multiples executions).

Performance on *test_2* partition

As with the *test_1* partition, the AUC values of the PR curves with *automatic* and *manual* labels on *test_2* were obtained (see Table 2). In these tables, similar values are observed with both labels. However, as in *test_1*, the order of the methods varied significantly (Friedman: $\chi^2(2) = 785.37$, $p < 2.2e^{-16}$). Similarly, significant differences were found with *automatics* labels (ANOVA: $F(5, 246) = 2097$, $p < 2e^{-16}$). Analogously, significant differences were found (ANOVA: $F(5, 246) = 1553$, $p < 2e^{-16}$) on *manual* labels.

⁴Source available at *removed for blind evaluation*

Table 1: AUC of the PR curves after 42 replications with *automatic* and *manual* labels on *test_1*.

<i>Automatic</i> labels		<i>Manual</i> labels	
Model	AUC	Model	AUC
BGWA	0.412 ± 0.026 ^a	BGWA	0.440 ± 0.023 ^a
CNN_ATT	0.194 ± 0.022 ^b	CNN_ATT	0.239 ± 0.031 ^b
CNN	0.193 ± 0.027 ^b	CNN	0.235 ± 0.027 ^c
RESIDE	0.191 ± 0.013 ^b	PCNN	0.209 ± 0.028 ^d
PCNN	0.158 ± 0.023 ^c	RESIDE	0.199 ± 0.020 ^d
PCNN_ATT	0.151 ± 0.025 ^d	PCNN_ATT	0.197 ± 0.029 ^d

^a differences with rest of methods***.

^b differences with BGWA***, PCNN*** and PCNN_ATT***.

^c differences with rest of methods*** except PCNN_ATT.)

^d differences with rest of methods*** except PCNN.)

*, **, *** to indicate $p < 0.05$, $p < 0.01$ and $p < 0.001$ respectively.

^a differences with rest of methods***.

^b differences with rest of methods*** except CNN.

^c differences with rest of methods*** except CNN_ATT.

^d differences with BGWA***, CNN*** and CNN_ATT***.

Table 2: AUC of the PR curves after 42 replications with *automatic* and *manual* labels on *test_2*.

<i>Automatic</i> labels		<i>Manual</i> labels	
Model	AUC	Model	AUC
BGWA	0.339 ± 0.016 ^a	BGWA	0.345 ± 0.021 ^a
PCNN_ATT	0.112 ± 0.015 ^b	PCNN_ATT	0.118 ± 0.017 ^b
CNN_ATT	0.105 ± 0.017 ^c	PCNN	0.109 ± 0.020 ^c
PCNN	0.105 ± 0.018 ^c	CNN_ATT	0.106 ± 0.018 ^d
CNN	0.098 ± 0.016 ^d	CNN	0.098 ± 0.017 ^e
RESIDE	0.021 ± 0.006 ^c	RESIDE	0.028 ± 0.011 ^f

^a differences with rest of methods***.

^b differences with BGWA*** and CNN***.

^c differences with BGWA***.

^d differences with BGWA*** and PCNN_ATT***.

*, **, *** to indicate $p < 0.05$, $p < 0.01$ and $p < 0.001$ respectively.

^a differences with rest of methods***.

^b differences with CNN*** and CNN_ATT*.

^c differences with BGWA*** and CNN*.

^d differences with BGWA*** and PCNN_ATT***.

^e differences with BGWA***, PCNN_ATT*** and PCNN*.

^f differences with BGWA***

The Figures 4a and 4b show the PR curves in one of the executions made with *automatic* and *manual* labels respectively on *test_2*.

5.2 Precision at N

Performance on *test_1* partition

The P@25 and P@50 subsets from the *test_1* partition were established, in addition to all the instances (P@All). Table 3 shows that the order of the models remains the same for the first three models by increasing the N, unlike the last three positions. The same happens with Table 4, in this case, the first two models are kept. The order of the models, as with the AUC, varied significantly for the *automatic* and *manual* labels on P@All (Friedman: $\chi^2(2) = 382.28$, $p < 2.2e^{-16}$). Similarly, there are significant differences in the performance of methods with *automatic* (ANOVA: $F(5, 246) = 210.8$, $p < 2e^{-16}$) and *manual* (ANOVA: $F(5, 246) = 255.6$, $p < 2e^{-16}$) labels. Then, two-by-two comparisons with Holm Correction (Holm, 1979) show significant differences with *automatic* labels

between the BGWA and RESIDE models and the rest. Similarly, two-by-two comparisons show significant differences with manual labels between the BGWA model and the rest. In addition, PCNN_ATT has significant differences with the other models except for PCNN (in reverse order it also happens). In this case, RESIDE only shows significant differences with BGWA, PCNN and PCNN_ATT.

Performance on *test_2* partition

In the same way as with *test_1*, the subsets P@25 and P@50 were established together with P@All, which includes the entire set. With both labeled, only two methods did not vary their order in the three subsets, BGWA and RESIDE (see Tables 5 and 6). On the other hand, the order of the methods using the P@All results varied significantly with respect to the *automatic* and *manual* labels (Friedman: $\chi^2(2) = 369.55$, $p < 2.2e^{-16}$)⁵.

⁵It should be noted that in all cases the Friedman test is used on the ranking of each execution, not only on the final results.

Table 3: P@25, P@50 and P@All after 42 replications with *automatic* labels on *test_1*.

Model	P@25	Model	P@50	Model	P@All
BGWA	0.819±0.062	BGWA	0.730±0.041	BGWA	0.558±0.029
CNN	0.587±0.087	CNN	0.489±0.062	CNN	0.386±0.036
CNN_ATT	0.580±0.089	CNN_ATT	0.486±0.064	CNN_ATT	0.375±0.045
PCNN	0.554±0.087	PCNN_ATT	0.461±0.055	PCNN	0.362±0.037
RESIDE	0.552±0.074	PCNN	0.459±0.060	PCNN_ATT	0.351±0.040
PCNN_ATT	0.550±0.079	RESIDE	0.433±0.054	RESIDE	0.325±0.035

Table 4: P@25, P@50 and P@All after 42 replications with *manual* labels on *test_1*.

Model	P@25	Model	P@50	Model	P@All
BGWA	0.715±0.079	BGWA	0.677±0.044	BGWA	0.585±0.033
RESIDE	0.555±0.075	RESIDE	0.489±0.043	RESIDE	0.376±0.037
CNN	0.551±0.089	CNN_ATT	0.465±0.061	CNN	0.370±0.035
CNN_ATT	0.544±0.089	CNN	0.459±0.059	CNN_ATT	0.370±0.044
PCNN	0.486±0.093	PCNN	0.401±0.062	PCNN	0.328±0.044
PCNN_ATT	0.458±0.096	PCNN_ATT	0.399±0.066	PCNN_ATT	0.325±0.041

Similarly, significant differences were found in the performance of the methods with *automatic* (ANOVA: $F(5, 246) = 1610, p < 2e^{-16}$) and *manual* (ANOVA: $F(5, 246) = 1265, p < 2e^{-16}$) labels. Then, in two-by-two comparisons with Holm Correction (Holm, 1979) there are no significant differences only between the CNN and CNN_ATT and PCNN and PCNN_ATT methods with both labeled.

6 Discussion

Our results indicate that the order of the algorithms, in terms of AUC of the PR curves on *test_1* and *test_2* partition, differ depending on the labeling. This justifies our claim that the interpretation of the PR curves, when used for evaluating DS algorithms, must be reconsidered. PR curves using automatic labels as a reference is not an optimal way to compare methods performance in DS, because it breaks a premise of the PR curves construction; that *true* labels are available. Several authors (Riedel, Yao, and McCallum, 2010; Hoffmann et al., 2011; Surdeanu et al., 2012; Zeng et al., 2015; Lin et al., 2016; Jiang et al., 2016; Liu et al., 2017; Vashishth et al., 2018; Ru et al., 2018; Zhou et al., 2018; Wang et al., 2018; Jat, Khandelwal, and Talukdar, 2018; Wu, Fan, and Zhang, 2019; Xu and Barbosa, 2019; Ye and Ling, 2019) have based the comparison of their method on the PR curves on these labels. The classical interpretation does not provide guarantees as to which method is performing better or which one is more tolerant

to noise in the labels.

On the other hand, the Section 5.2 has also confirmed that P@N is not being interpreted correctly in DS either. This is critical for the task at hand considering the unbalance in the data sets, variability among the relations, selection criteria, among others. There is not a clear criterion or criteria that guarantee to choose the same instances for the evaluation of each of the methods. In other words, it is not guaranteed that the first instances chosen to evaluate one method are the same for another method. If the selection is based on the score of the classifier, it varies from one to another execution. The same happens if the selection is random. For example, it may happen that for a method the first N instances are of the same relation. This indicates how good this method is for that relation, however, for the rest, its performance is not known. Also, sometimes, the P@N is calculated over *automatic* labels whereas some works do it over *manual* labels. This is the case of the 6 methods used in this work. This further confuses interpretation. Furthermore, dispersion values are not reported in the aforementioned works which mathematically renders those works uninformative.

What was expressed above shows that PR curves and P@N measures are not currently being interpreted properly in DS due to the presence of noisy labels. Currently, we believe there are no reliable statistics regarding the actual performance of the DS meth-

Table 5: P@25, P@50 and P@All after 42 replications with *automatic* labels on *test_2*.

Model	P@25	Model	P@50	Model	P@All
BGWA	0.804 ± 0.082	BGWA	0.762 ± 0.064	BGWA	0.019 ± 0.000
CNNATT	0.360 ± 0.112	CNN	0.357 ± 0.084	CNN	0.015 ± 0.000
CNN	0.346 ± 0.111	CNNATT	0.341 ± 0.087	CNNATT	0.015 ± 0.000
PCNNATT	0.273 ± 0.089	PCNNATT	0.268 ± 0.067	PCNN	0.014 ± 0.000
PCNN	0.252 ± 0.106	PCNN	0.233 ± 0.070	PCNNATT	0.014 ± 0.000
RESIDE	0.115 ± 0.095	RESIDE	0.129 ± 0.076	RESIDE	0.010 ± 0.000

Table 6: P@25, P@50 and P@All after 42 replications with *manual* labels on *test_2*.

Model	P@25	Model	P@50	Model	P@All
BGWA	0.017 ± 0.000	BGWA	0.795 ± 0.083	BGWA	0.0168 ± 0.000
CNN	0.014 ± 0.000	CNNATT	0.343 ± 0.120	CNN	0.0137 ± 0.000
CNNATT	0.014 ± 0.000	CNN	0.320 ± 0.117	CNNATT	0.0135 ± 0.000
PCNN	0.013 ± 0.000	PCNNATT	0.255 ± 0.104	PCNN	0.0130 ± 0.000
PCNNATT	0.013 ± 0.000	PCNN	0.230 ± 0.099	PCNNATT	0.0129 ± 0.000
RESIDE	0.010 ± 0.000	RESIDE	0.150 ± 0.094	RESIDE	0.0103 ± 0.000

ods. While the community agrees on a mathematically correct interpretation in this context, or new statistics are proposed for evaluating the performance of DS methods, a possible strategy to circumvent the deadlock is what was done here. That is, selecting multiple instances of the evaluation data set while maintaining its distribution (*test_1* partition). Then, perform a manual review of these instances using multiple raters. The main limitations of *test_1* partition are the instances number selected. This is why the *test_2* partition was labeled with multiple raters using MTurk. The advantage of this partition with respect to *test_1* and those proposed by (Hoffmann et al., 2011) and (Ren et al., 2017) is that it is made up of all the instances of the NYT2010 data set test partition (only those different from NA were labeled with MTurk). From the *test_2* partition, the methods can be compared with precision, recall and F1 using the traditional interpretation.

Gasta un par de frase o 1 párrafo en validación nomológica. Limitaciones?

7 Conclusions

Significant differences were found in the ordering of the methods regarding their performances, when the performance is established according to the AUC of the PR curves between the evaluation using the *automatic* labels and the same data set with the *manual* labels. The largest AUCs were obtained using *manual* labels which speaks well of the

capacity of the DS methods to handle noisy data as it is their core intention.


Our results suggest that PR curves are currently not being interpreted correctly in DS and, on the other hand, manual evaluation of the first N instances does not cover the entire data set. The existing selection criteria for the instances to be manually reviewed are not deterministic, suggesting multiple executions of the method and the dispersion report. Besides, these measures, as they are being used, are inconclusive as to the performance of those methods. **Estamos repitiendo los resultados. Las conclusiones deben ir más allá y explorar las implicaciones.** Finally, a partition was provided that allows you to evaluate this task using labels manually reviewed by multiple raters. This partition also allows the use of precision, recall and F1 measurements. **Trabajo futuro?**

Acknowledgments


removed for blind evaluation

References

- Agichtein, E. and L. Gravano. 2000. Snowball: Extracting Relations from large Plain-Text Collections. In *Proceedings of the fifth ACM conference on Digital libraries*, pages 85–94. ACM.
- Bastos, A., A. Nadgeri, K. Singh, I. O. Mulang, S. Shekarpour, J. Hoffart, and M. Kaul. 2020. RECON: Relation Extraction using Knowledge Graph Context



(a) Automatic labels




(b) Manual labels

Figure 3: PR curves corresponding to evaluation of the DS algorithms over *test_1* (one execution) set pick for verification in (a) *automatic* labels and (b) *manual* labels. The AUC of the PR curves is indicated beside each label in the legend.

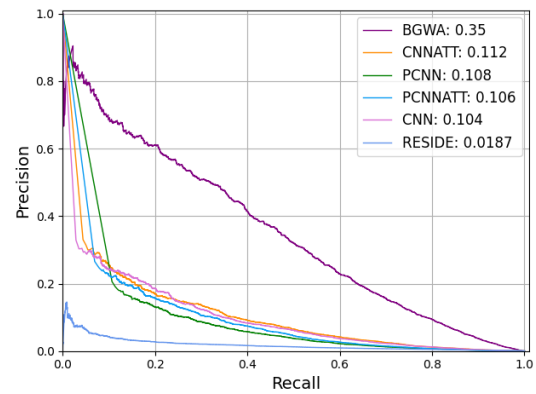
in a Graph Neural Network. In *Proceedings of the Web Conference 2021*, pages 1673–1685, Ljubljana.

Bunescu, R. C. and R. J. Mooney. 2005. A Shortest Path Dependency Kernel for Relation Extraction. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP)*, pages 724–731, Vancouver, Association for Computational Linguistics.

Davis, J. and M. Goadrich. 2006. The relationship between Precision-Recall and ROC curves. In *Proceedings of the 23rd International Conference on Machine Learning - ICML '06*, pages 233–240, Pittsburgh, Pennsylvania, USA. ACM Press.



(a) Automatic labels



(b) Manual labels

Figure 4: PR curves corresponding to evaluation of the DS algorithms over *test_2* (one execution) set pick for verification in (a) *automatic* labels and (b) *manual* labels. The AUC of the PR curves is indicated beside each label in the legend.

Defferrard, M., X. Bresson, and P. Vandergheynst. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *Advances in neural information processing systems*, pages 3844–3852.

Friedman, M. 1940. A comparison of alternative tests of significance for the problem of m rankings. *The Annals of Mathematical Statistics*, 11(1):86–92.

He, Z., W. Chen, Z. Li, M. Zhang, W. Zhang, and M. Zhang. 2018. SEE: Syntax-Aware Entity Embedding for Neural Relation Extraction. In *The Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 5795–5802. Association for the Advancement of Artificial Intelligence.

Hearst, M. A. 1992. Automatic Acquisition

- of Hyponyms from Large Text Corpora. In *Proceedings of the 14th Conference on Computational Linguistics (COLING-92)*, pages 539–545, Nantes.
- Hoffmann, R., C. Zhang, X. Ling, L. Zettlemoyer, and D. S. Weld. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 541–550, Portland, Oregon. Association for Computational Linguistics.
- Holm, S. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 6(2):65–70.
- Howell, D. C. 2012. *Statistical Methods for Psychology*. Cengage Learning ALL, eighth edi edition.
- Jat, S., S. Khandelwal, and P. Talukdar. 2018. Improving Distantly Supervised Relation Extraction using Word and Entity Based Attention. *arXiv:1804.06987v1 [cs.CL]*, apr.
- Ji, G., K. Liu, S. He, and J. Zhao. 2017. Distant Supervision for Relation Extraction with Sentence-Level Attention and Entity Descriptions. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*, pages 3060–3066.
- Jia, W., D. Dai, X. Xiao, and H. Wu. 2019. ARNOR: Attention Regularization based Noise Reduction for Distant Supervision Relation Classification. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1399–1408, Florence. Association for Computational Linguistics.
- Jiang, X., Q. Wang, P. Li, and B. Wang. 2016. Relation extraction with multi-instance multi-label convolutional neural networks. In *Proceedings of the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1471–1480, Osaka.
- Keilwagen, J., I. Grosse, and J. Grau. 2014. Area under Precision-Recall Curves for Weighted and Unweighted Data. *PLoS ONE*, 9(3):e92209, mar.
- Lin, Y., S. Shen, Z. Liu, H. Luan, and M. Sun. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 2124–2133, Berlin, Germany. Association for Computational Linguistics.
- Liu, T., K. Wang, B. Chang, and Z. Sui. 2017. A Soft-label Method for Noise-tolerant Distantly Supervised Relation Extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1790–1795, Copenhagen, Denmark.
- Manning, C. D., P. Raghavan, and H. Schütze. 2008. *Introduction to information retrieval*. Cambridge University Press.
- Mintz, M., S. Bills, R. Snow, and D. Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the 47th Annual Meeting of the ACL*, pages 1003–1011, Suntec, Singapore.
- Piskorski, J. and R. Yangarber. 2013. Information extraction: Past, Present and Future. In *Multi-source, Multilingual Information Extraction and Summarization 11*. Springer-Verlag Berlin Heidelberg, pages 23–49.
- Ren, X., Z. Wu, W. He, M. Qu, C. R. Voss, H. Ji, T. F. Abdelzaher, and J. Han. 2017. CoType. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1015–1024, Perth, apr. International World Wide Web Conferences Steering Committee.
- Riedel, S., L. Yao, and A. McCallum. 2010. Modeling relations and their mentions without labeled text. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 148–163, Berlin. Springer.
- Ru, C., J. Tang, S. Li, S. Xie, and T. Wang. 2018. Using semantic similarity to reduce wrong labels in distant supervision for relation extraction. *Information Processing Management*, 54(4):593–608, jul.
- Student. 1908. The Probable Error of a Mean. *Biometrika*, 6(1):1–25.

- Surdeanu, M., J. Tibshirani, R. Nallapati, and C. D. Manning. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 455–465, Jeju Island, Korea. Association for Computational Linguistics.
- Vashishth, S., R. Joshi, S. S. Prayaga, C. Bhattacharyya, and P. Talukdar. 2018. Reside: Improving Distantly-Supervised Neural Relation Extraction using Side Information. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1257–1266, Brussels, Belgium. Association for Computational Linguistics.
- Wang, G., W. Zhang, R. Wang, Y. Zhou, L. Chen, W. Zhang, H. Zhu, and H. Chen. 2018. Label-free distant supervision for relation extraction via knowledge graph embedding. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 2246–2255.
- Wu, S., K. Fan, and Q. Zhang. 2019. Improving Distantly Supervised Relation Extraction with Neural Noise Converter and Conditional Optimal Selector. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7273–7280, nov.
- Xu, P. and D. Barbosa. 2019. Connecting Language and Knowledge with Heterogeneous Representations for Neural Relation Extraction. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3201–3206, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ye, Z.-X. and Z.-H. Ling. 2019. Distant Supervision Relation Extraction with Intra-Bag and Inter-Bag Attentions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2810–2819, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zeng, D., K. Liu, Y. Chen, and J. Zhao. 2015. Distant Supervision for Relation Extraction via Piecewise Convolutional Neural Networks. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1753–1762, Lisbon, Portugal. Association for Computational Linguistics.
- Zeng, D., K. Liu, S. Lai, G. Zhou, and J. Zhao. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344, Dublin, Ireland.
- Zhou, P., J. Xu, Z. Qi, H. Bao, Z. Chen, and B. Xu. 2018. Distant supervision for relation extraction with hierarchical selective attention. *Neural Networks*, 108:240–247.