On the Evaluation and Comparison of Taggers: the Effect of Noise in Testing Corpora.

Lluís Padró and Lluís Màrquez
Dep. LSI. Technical University of Catalonia
c/ Jordi Girona 1-3. 08034 Barcelona
{padro,lluism}@lsi.upc.es

Abstract

This paper addresses the issue of POS tagger evaluation. Such evaluation is usually performed by comparing the tagger output with a reference test corpus, which is assumed to be error-free. Currently used corpora contain noise which causes the obtained performance to be a distortion of the real value. We analyze to what extent this distortion may invalidate the comparison between taggers or the measure of the improvement given by a new system. The main conclusion is that a more rigorous testing experimentation setting/designing is needed to reliably evaluate and compare tagger accuracies.

1 Introduction and Motivation

Part of Speech (POS) Tagging is a quite well defined NLP problem, which consists of assigning to each word in a text the proper morphosyntactic tag for the given context. Although many words are ambiguous regarding their POS, in most cases they can be completely disambiguated taking into account an adequate context. Successful taggers have been built using several approaches, such as statistical techniques, symbolic machine learning techniques, neural networks, etc. The accuracy reported by most current taggers ranges from 96–97% to almost 100% in the linguistically-motivated Constraint Grammar environment.

Unfortunately, there have been very few direct comparisons of alternative taggers¹ on identical test data. However, in most current papers it is argued that the performance of some taggers is better than others as a result of some kind of indirect comparisons between them. We

think that there are a number of not enough controlled/considered factors that make these conclusions dubious in most cases.

In this direction, the present paper aims to point out some of the difficulties arising when evaluating and comparing tagger performances against a reference test corpus, and to make some criticism about common practices followed by the NLP researchers in this issue.

The above mentioned factors can affect either the evaluation or the comparison process. Factors affecting the evaluation process are: (1) Training and test experiments are usually performed over noisy corpora which distorts the obtained results, (2) performance figures are too often calculated from only a single or very small number of trials, though average results from multiple trials are crucial to obtain reliable estimations of accuracy (Mooney, 1996), (3) testing experiments are usually done on corpora with the same characteristics as the training data –usually a small fresh portion of the training corpus– but no serious attempts have been done in order to determine the reliability of the results when moving from one domain to another (Krovetz, 1997), and (4) no figures about computational effort –space/time complexity– are usually reported, even from an empirical perspective. A factors affecting the comparison process is that comparisons between taggers are often indirect, while they should be compared under the same conditions in a multiple-trial experiment with statistical tests of significance.

For these reasons, this paper calls for a discussion on POS taggers evaluation, aiming to establish a more rigorous test experimentation setting/designing, indispensable to extract reliable conclusions. As a starting point, we will focus only on how the noise in the test corpus can affect the obtained results.

¹One of the exceptions is the work by (Samuelsson and Voutilainen, 1997), in which a very strict comparison between taggers is performed.

2 Noise in the testing corpus

From a machine learning perspective, the relevant noise in the corpus is that of non systematically mistagged words (i.e. different annotations for words appearing in the same syntactic/semantic contexts).

Commonly used annotated corpora have noise. See, for instance, the following examples from the Wall Street Journal (WSJ) corpus:

Verb participle forms are sometimes tagged as such (VBN) and also as adjectives (JJ) in other sentences with no structural differences:

- 1a) ... failing_VBG to_TO voluntarily_RB
submit_VB the_DT requested_VBN
information_NN ...
- 1b) ... a_DT large_JJ sample_NN of_IN
married_JJ women_NNS with_IN at_IN
least_JJS one_CD child_NN ...

Another structure not coherently tagged are noun chains when the nouns (NN) are ambiguous and can be also adjectives (JJ):

- 2a) ... Mr._NNP Hahn_NNP ,-, the_DT
62-year-old_JJ chairman_NN and_CC
chief_NN executive_JJ officer_NN of_IN
Georgia-Pacific_NNP Corp._NNP ...
- 2b) ... Burger_NNP King_NNP
's_POS chief_JJ executive_NN officer_NN ,-,
Barry_NNP Gibbons_NNP ,-, stars_VBZ
in_IN ads_NNS saying_VBG ...

The noise in the test set produces a wrong estimation of accuracy, since correct answers are computed as wrong and vice-versa. In following sections we will show how this uncertainty in the evaluation may be, in some cases, larger than the reported improvements from one system to another, so invalidating the conclusions of the comparison.

3 Model Setting

To study the appropriateness of the choices made by a POS tagger, a reference tagging must be selected and assumed to be correct in order to compare it with the tagger output. This is usually done by assuming that the disambiguated test corpora being used contains the right POS disambiguation. This approach is quite right when the tagger error rate is larger enough than the test corpus error rate, nevertheless, the current POS taggers have reached a performance level that invalidates this choice,

since the tagger error rate is getting too close to the error rate of the test corpus.

Since we want to study the relationship between the tagger error rate and the test corpus error rate, we have to establish an absolute reference point. Although (Church, 1992) questions the concept of *correct analysis*, (Samuelsson and Voutilainen, 1997) establish that there exists a –statistically significant– *absolute* correct disambiguation, respect to which the error rates of either the tagger or the test corpus can be computed. What we will focus on is how distorted is the tagger error rate by the use of a noisy test corpus as a reference.

The cases we can find when evaluating the performance of a certain tagger are presented in table 1. OK/¬OK stand for a right/wrong tag (respect to the absolute correct disambiguation). When both the tagger and the test corpus have the correct tag, the tag is correctly evaluated as *right*. When the test corpus has the correct tag and the tagger gets it wrong, the occurrence is correctly evaluated as *wrong*. But problems arise when the test corpus has a wrong tag: If the tagger gets it correctly, it is evaluated as *wrong* when it should be *right* (false negative). If the tagger gets it wrong, it will be rightly evaluated as *wrong* if the error committed by the tagger is other than the error in the test corpus, but wrongly evaluated as *right* (false positive) if the error is the same. Table 1 shows the computation of the percent-

corpus	tagger	eval: <i>right</i>	eval: <i>wrong</i>
OK _c	OK _t	$(1-C)t$	–
OK _c	¬OK _t	–	$(1-C)(1-t)$
¬OK _c	OK _t	–	Cu
¬OK _c	¬OK _t	$C(1-u)p$	$C(1-u)(1-p)$

Table 1: Possible cases when evaluating a tagger.

ages of each case. The meanings of the used variables are:

- C*: Test corpus error rate. Usually an estimation is supplied with the corpus.
- t*: Tagger performance rate on words rightly tagged in the test corpus. It can be seen as $P(\text{OK}_t | \text{OK}_c)$.
- u*: Tagger performance rate on words wrongly tagged in the test corpus. It can be seen as $P(\text{OK}_t | \neg \text{OK}_c)$.

p : Probability that the tagger makes the same error as the test corpus, given that both get a wrong tag.

x : Real performance of the tagger, i.e. what would be obtained on an error-free test set.

K : Observed performance of the tagger, computed on the noisy test corpus.

For simplicity, we will consider only performance on ambiguous words. Considering unambiguous words will make the analysis more complex, since it should be taken into account that neither the behaviour of the tagger (given by u , t , p) nor the errors in the test corpus (given by c) are the same on ambiguous and unambiguous words. Nevertheless, this is an issue that must be further addressed.

If we knew each one of the above proportions, we would be able to compute the real performance of our tagger (x) by adding up the OK_t rows from table 1, i.e. the cases in which the tagger got the right disambiguation independently from the tagging of the test set:

$$x = (1-C)t + Cu \quad (1)$$

The equation of the observed performance can also be extracted from table 1, adding up what is evaluated as *right*:

$$K = (1-C)t + C(1-u)p \quad (2)$$

The relationship between the real and the observed performance is derived from 1 and 2:

$$x = K - C(1-u)p + Cu$$

Since only K and C are known (or approximately estimated) we can not compute the real performance of the tagger. All we can do is to establish some reasonable bounds for t , u and p , and see in which range is x .

Since all variables are probabilities, they are bounded in $[0, 1]$. We also can assume² that $K > C$. We can use this constraints and the above equations to bound the values of all variables. From 2, we obtain:

$$u = 1 - \frac{K - t(1-C)}{Cp}; \quad p = \frac{K - t(1-C)}{C(1-u)}; \quad t = \frac{K - C(1-u)p}{1-C}$$

Thus, u will be maximum when p and t are maximum (i.e. 1). This gives an upper bound

²In the cases we are interested in –that is, current systems– the tagger observed performance, K , is over 90%, while the corpus error rate, C , is below 10%.

for u of $(1-K)/C$. When $t = 0$, u will range in $[-\infty, 1-K/C]$ depending on the value of p . Since we are assuming $K > C$, the most informative lower bound for u keeps being zero. Similarly, p is minimum when $t=1$ and $u=0$. When $t = 0$ the value for p will range in $[K/C, +\infty]$ depending on u . Since $K > C$, the most informative upper bound for p is still 1. Finally, t will be maximum when $u = 1$ and $p = 0$, and minimum when $u=0$ and $p=1$. Summarizing:

$$0 \leq u \leq \min \left\{ 1, \frac{1-K}{C} \right\} \quad (3)$$

$$\max \left\{ 0, \frac{K+C-1}{C} \right\} \leq p \leq 1 \quad (4)$$

$$\frac{K-C}{1-C} \leq t \leq \min \left\{ 1, \frac{K}{1-C} \right\} \quad (5)$$

Since the values of the variables are mutually constrained, it is not possible that, for instance, u and t have simultaneously their upper bound values (if $(1-K)/C < 1$ then $K/(1-C) > 1$ and viceversa). Any bound which is out of $[0, 1]$ is not informative and the appropriate boundary, 0 or 1, is then used. Note that the lower bound for t will never be negative under the assumption $K > C$.

Once we have established these bounds, we can use equation 1 to compute the range for the real performance value of our tagger: x will be minimum when u and t are minimum, which produces the following bounds:

$$x_{min} = K - Cp \quad (6)$$

$$x_{max} = \begin{cases} K+C & \text{if } K \leq 1-C \\ 1 - \frac{K+C-1}{p} & \text{if } K \geq 1-C \end{cases} \quad (7)$$

As an example, let's suppose we evaluate a tagger on a test corpus which is known to contain about 3% of errors ($C=0.03$), and obtain a reported performance of 93%³ ($K=0.93$). In this case, equations 6 and 7 yield a range for the real performance x that varies from $[0.93, 0.96]$ when $p=0$ to $[0.90, 0.96]$ when $p=1$.

This results suggest that although we observe a performance of K , we can not be sure of how well is our tagger performing without taking into account the values of t , u and p .

It is also obvious that the intervals in the above example are too wide, since they consider all the possible parameter values, even when they correspond to very unlikely param-

³This is a realistic case obtained by (Màrquez and Padró, 1997) tagger. Note that 93% is the accuracy on ambiguous words (the equivalent overall accuracy was about 97%).

eter combinations⁴. In section 4 we will try to narrow those intervals, limiting the possibilities to *reasonable* cases.

4 Reasonable Bounds for the Basic Parameters

In real cases, not all parameter combinations will be equally likely. In addition, the bounds for the values of t , u and p are closely related to the similarities between the training and test corpora. That is, if the training and test sets are extracted from the same corpus, they will probably contain the same kind of errors in the same kind of situations. This may cause the training procedure to *learn* the errors—especially if they are systematic—and thus the resulting tagger will tend to make the same errors that appear in the test set. On the contrary, if the training and test sets come from different sources—sharing only the tag set—the behaviour of the resulting tagger will not depend on the right or wrong tagging of the test set.

We can try to establish narrower bounds for the parameters than those obtained in section 3.

First of all, the value of t is already constrained enough, due to its high contribution $(1 - C)$ to the value of K , which forces t to take a value close to K . For instance, applying the boundaries in equation 5 to the case $C=0.03$ and $K=0.93$, we obtain that t belongs to $[0.928, 0.959]$.

The range for u can be slightly narrowed considering the following: In the case of independent test and training corpora, u will tend to be equal to t . Otherwise, the more biased towards the corpus errors is the language model, the lower u will be. Note that $u > t$ would mean that the tagger disambiguates *better* the noisy cases than the correct ones. Concerning to the lower bound, only in the case that all the errors in the training and test corpus were systematic (and thus can be learned) could u reach zero. However, not only this is not a likely situation, but also requires a perfect-learning tagger. It seems more reasonable that, in normal cases, errors will be random, and the tagger will behave

randomly on the noisy occurrences. This yields a lower bound for u of $1/a$, being a the average ambiguity ratio for ambiguous words.

The *reasonable* bounds for u are thus

$$\frac{1}{a} \leq u \leq \min \left\{ t, \frac{1-K}{C} \right\}$$

Finally, the value of p has similar constraints to those of u . If the test and training corpora are independent, the probability of making the same error, given that both are wrong, will be the random $1/(a-1)$. If the corpora are not independent, the errors that can be learned by the tagger will cause p to rise up to (potentially) 1. Again, only in the case that all errors were systematic, could p reach 1.

Then, the *reasonable* bounds for p are:

$$\max \left\{ \frac{1}{a-1}, \frac{K+C-1}{C} \right\} \leq p \leq 1$$

5 On Comparing Tagger Performances

As stated above, knowing which are the *reasonable* limits for the u , p and t parameters enables us to compute the range in which the real performance of the tagger can vary.

So, given two different taggers T_1 and T_2 , and provided we know the values for the test corpus error rate and the observed performance of both cases (C_1, C_2, K_1, K_2) , we can compare them by matching the *reasonable* intervals for the respective real performances x_1 and x_2 .

From a conservative position, we cannot strongly state than one of the taggers performs better than the other when the two intervals overlap, since this implies a chance that the real performances of both taggers are the same.

The following real example has been extracted from (Màrquez and Padró, 1997): The tagger T_1 uses only bigram information and has an observed performance on ambiguous words $K_1 = 0.9135$ (96.86% overall). The tagger T_2 uses trigrams and automatically acquired context constraints and has an accuracy of $K_2 = 0.9282$ (97.39% overall). Both taggers have been evaluated on a corpus (WSJ) with an estimated error rate⁵ $C_1 = C_2 = 0.03$. The average ambiguity ratio of the ambiguous words in the corpus is $a = 2.5$ tags/word.

⁴For instance, it is not reasonable that $u=0$, which would mean that the tagger **never** disambiguates correctly a wrong word in the corpus, or $p=1$, which would mean that it **always** makes the same error when both are wrong.

⁵The (WSJ) corpus error rate is estimated over all words. We are assuming that the errors distribute uniformly among all words, although ambiguous words

These data yield the following range of reasonable intervals for the real performance of the taggers.

for $p_i = (1/a) = 0.4$	for $p_i = 1$
$x_1 \in [91.35, 94.05]$	$x_1 \in [90.75, 93.99]$
$x_2 \in [92.82, 95.60]$	$x_2 \in [92.22, 95.55]$

The same information is included in figure 1 which presents the reasonable accuracy intervals for both taggers, for p ranging from $1/a = 0.4$ to 1 (the shadowed part corresponds to the overlapping region between intervals).

Figure 1: Reasonable intervals for both taggers

The obtained intervals have a large overlap region which implies that there are *reasonable* parameter combinations that could cause the taggers to produce different observed performances though their real accuracies were very similar. From this conservative approach, we would not be able to conclude that the tagger T_2 is better than T_1 , even though the 95% confidence intervals for the observed performances did allow us to do so.

6 Discussion

The presented analysis of the effects of noise in the test corpus on the evaluation of POS taggers leads us to conclude that when a tagger is evaluated as better than another using noisy test corpus, there are *reasonable* chances that they are in fact very similar but one of them is just adapting better than the other to the noise in the corpus.

probably have a higher error rate. Nevertheless, a higher value for C would cause the intervals to be wider and to overlap even more.

We believe that the widespread practice of evaluating taggers against a noisy test corpus has reached its limit, since the performance of current taggers is getting too close to the error rate usually found in test corpora.

An obvious solution –and maybe not as costly as one might think, since small test sets properly used may yield enough statistical evidence– is using only error-free test corpora. Another possibility is to further study the influence of noise in order to establish a criterion –e.g. a threshold depending on the amount of overlapping between intervals– to decide whether a given tagger can be considered better than another.

There is still much to be done in this direction. This paper does not intend to establish a new evaluation method for POS tagging, but to point out that there are some issues –such as the noise in test corpus– that have been paid little attention and are more important than what they seem to be.

Some of the issues that should be further considered are: The effect of noise on unambiguous words; the reasonable intervals for *overall* real performance; the –probably– different values of C , p , u and t for ambiguous/unambiguous words; how to estimate the parameter values of the evaluated tagger in order to constrain as much as possible the intervals; the statistical significance of the interval overlappings; a more informed (and less conservative) criterion to reject/accept the hypothesis that both taggers are different, etc.

References

- Church, K.W. 1992. Current Practice in Part of Speech Tagging and Suggestions for the Future. In Simmons (ed.), *Sbornik praci: In Honor of Henry Kučera*. Michigan Slavic Studies.
- Krovetz, R. 1997. Homonymy and Polysemy in Information Retrieval. In *Proceedings of joint E/ACL meeting*.
- Màrquez, L. and Padró, L. 1997. A Flexible POS Tagger Using an Automatically Acquired Language Model. In *Proceedings of joint E/ACL meeting*.
- Mooney, R.J. 1996. Comparative Experiments on Disambiguating Word Senses: An Illustration of the Role of Bias in Machine Learning. In *Proceedings of EMNLP'96 conference*.
- Samuelsson, C. and Voutilainen, A. 1997. Comparing a Linguistic and a Stochastic Tagger. In *Proceedings of joint E/ACL meeting*.

Resum

Aquest article versa sobre l'avaluació de desambiguadors morfosintàctics. Normalment, l'avaluació es fa comparant la sortida del desambiguador amb un corpus de referència, que se suposa lliure d'errors. De tota manera, els corpus que s'usen habitualment contenen soroll que causa que el rendiment que s'obté dels desambiguadors sigui una distorsió del valor real. En aquest article analitzem fins a quin punt aquesta distorsió pot invalidar la comparació entre desambiguadors o la mesura de la millora aportada per un nou sistema. La conclusió principal és que cal establir procediments alternatius d'experimentació més rigorosos, per poder avaluar i comparar fiablement les precisions dels desambiguadors morfosintàctics.

Laburtena

Artikulu hau desanbiguatzaile morfosintaktikoen ebaluazioaren inguruan datza. Normalean, ebaluazioa, desanbiguatzailearen irteera eta ustez errorerik gabeko erreferentziako corpus bat konparatuz egiten da. Hala ere, maiz corpusetan erroreak egoten dira eta horrek desanbiguatzailearen emaitzaren benetako balioan eragina izaten du. Artikulu honetan, hain zuzen ere, horixe aztertuko dugu, alegia, zer neurritan distortsio horrek jar dezakeen auzitan desanbiguatzaileen arteko konparazioa edo sistema berri batek ekar dezakeen hobekuntza-maila. Konklusiorik nagusiena hau da: desanbiguatzaile morfosintaktikoak aztertzeke eta modu ziurrago batez konparatu ahal izateko, azterketa-bideak sakonagoak eta zehatzagoak izan beharko liratekeela.