

# DECOUR: a corpus of DEceptive statements in Italian COURts

Tommaso Fornaciari<sup>2</sup>, Massimo Poesio<sup>1,2</sup>

<sup>1</sup> University of Essex, Language and Computation Group

<sup>2</sup> University of Trento, Center for Mind/Brain Sciences

tommaso.fornaciari@unitn.it, massimo.poesio@unitn.it

## Abstract

In criminal proceedings, sometimes it is not easy to evaluate the sincerity of oral testimonies. DECOUR - DEception in COURt corpus - has been built with the aim of training models suitable to discriminate, from a stylometric point of view, between sincere and deceptive statements. DECOUR is a collection of hearings held in four Italian Courts, in which the speakers lie in front of the judge. These hearings become the object of a specific criminal proceeding for calumny or false testimony, in which the deceptiveness of the statements of the defendant is ascertained. Thanks to the final Court judgment, that points out which lies are told, each utterance of the corpus has been annotated as true, uncertain or false, according to its degree of truthfulness. Since the judgment of deceptiveness follows a judicial inquiry, the annotation has been realized with a greater degree of confidence than ever before. Moreover, in Italy this is the first corpus of deceptive texts not relying on ‘mock’ lies created in laboratory conditions, but which has been collected in a natural environment.

**Keywords:** Detecting deception, forensic linguistics, stylometry.

## 1. Introduction

Being able to detect deception in human interaction would be useful for a number of applications in which important decisions depend on the truthfulness of communication. However, finding reliable clues of deceptiveness is difficult. Experimental evidence suggests that human performance in recognizing deception is not much better than chance (Bond and De Paulo, 2006). Furthermore, in some studies human skills seem not to be particularly improved even after specific training (Levine et al., 2005).

Regarding deceptive communication there are many studies in literature in which verbal and non-verbal cues are taken into consideration. These studies can be divided into two main families, depending on how they solve the dilemma of data collection: in real life or in laboratory conditions. Both families of study present advantages and disadvantages. Laboratory studies allow control of experimental variables but, from a psychological point of view, they focus on ‘mock lies’: therefore their findings may not be extended to deception produced in real life. Instead, in field studies psychological conditions of subjects are genuine, but serious problems have to be faced. For example, in cases where linguistic clues are considered, it is often difficult - and sometimes impossible:

- to collect data in standardized conditions;
- to set up balanced data set;
- to verify sincerity of the statements.

For these reasons, linguists often studied deceptiveness in laboratory conditions. This is the case of Newman et al. (2003) and Strapparava and Mihalcea (2009), who collected corpora of sincere and deceptive texts coming from written and spoken language, and analyzed them with interesting results. The same approach has been employed in the context of Computer-Mediated Communication (Hancock et al., 2008; Zhou et al., 2004).

This focus on data collected in the laboratory is what led Lina Zhou to complain about ‘the lack of data sets for evaluating deception detection models’ (Zhou et al., 2008). In fact, one of the few examples of corpus collected in a natural environment regarding English language is represented by Fitzpatrick’s study (Bachenko et al., 2008; Fitzpatrick and Bachenko, 2009). Our goal was to produce a resource to study deceptive language (for Italian) overcoming the limitations said above.

In this paper we present the DECOUR corpus of deception in testimonies in Italian courts. In Section 2 we show how data have been collected; in Section 3 we describe how DECOUR corpus has been annotated; in Section 4 we summarize the characteristics of DECOUR corpus and in Section 5 the results of the preliminary analyses are discussed.

## 2. Data collection

In order to study deceptive language, we tried to build a corpus of texts:

- coming from a real life scenario;
- characterized by a strong psychological involvement of the speakers;
- collected in standard conditions;
- of which the truthfulness or truthlessness was known.

We found a way to fulfill these requirements in a legal context.

It happens in criminal proceedings that investigators interview, more or less formally, several subjects who consequently have the possibility to issue true or false statements. In most cases the reports in which the testimonies are collected do not bring back the words exactly pronounced by the subjects, but represent a synthesis of their declarations, carried out by the police officer who hears and records

them. These reports are not a faithful mirror of the linguistic behavior of the subjects, therefore they are not useful from the point of view of the present work.

In some particularly serious cases, it is also possible that the interrogation in front of the public prosecutor is recorded and transcribed word by word. These interrogations could be useful, but they are relatively rare and also difficult to find because in the proceedings where they could be carried out, they are not always. Above all, even when theoretically possible, to find external and objective evidences of the truthfulness or deceptiveness of statements would be very difficult from a practical point of view. In fact, these evidences are usually dispersed in a lot of different and various investigative data, often in a huge amount.

Therefore the point was to find testimonies not only recorded word by word, but also of which the truthfulness or deceptiveness was easily verifiable.

## 2.1. Hearings

In Italian criminal proceedings there is a specific moment in which all the testimonies are imperatively recorded word by word: that is, the hearings that take place during the debate in front of the judge. Furthermore, in some proceedings the truthfulness or deceptiveness of the testimonies is easily verifiable. It is the case of criminal proceedings ex art. 368 and 372 of the Italian Criminal Code, that concerns the crimes of ‘calumny’ and ‘false testimony’<sup>1</sup>. While the concept of false testimony is intuitive, in Italian Criminal Code calumny is a particular kind of false testimony, consisting of the attempt to charge on someone else the responsibility of a crime that has been committed. The distinction makes sense because in the Italian legal system nobody can be forced to say some truth unfavorable to oneself. It means that to lie about a committed crime is not a crime, but it is so if trying to charge the responsibility to someone else. In order to collect this kind of data, contacts have been taken with Courts in four Italian towns, with the aim to be allowed to examine their dossiers and extract information with scientific purposes. Authorizations have been received to collect data, with the only restriction of using them in anonymous form, in respect to the privacy of the involved subjects.

The inquiries for calumny and false testimony usually originate from another previous proceeding, in which the defendant or a witness takes part in a hearing and issues statements that are found not reliable. In these cases, a new

---

<sup>1</sup>The art. 368 reads: “Chiunque, con denuncia, querela, richiesta o istanza, anche se anonima o sotto falso nome, diretta all’Autorità giudiziaria o ad altra Autorità che a quella abbia obbligo di riferirne, incolpa di un reato taluno che egli sa innocente, ovvero simula a carico di lui le tracce di un reato, è punito con la reclusione da due a sei anni.” In brief, it punishes whoever tries to charge the responsibility of a crime on someone who he knows is innocent.

Art. 372 instead reads: “Chiunque, deponendo come testimone innanzi all’Autorità giudiziaria, afferma il falso o nega il vero, ovvero tace, in tutto o in parte ciò che sa intorno ai fatti sui quali è interrogato, è punito con la reclusione da due a sei anni.” This article punishes someone who, in front of the Judicial Authority, says a falsehood or denies the truth, or does not reveal what he knows about the investigated facts.

criminal proceeding arises, aimed to establish if the subject committed the crime of calumny or of false testimony. More rarely, the proceeding concerns statements which are not issued in a hearing, but in circumstances in which the words of the subjects are not recorded *verbatim*: typically, this is the case of the complaints lodged to the police. Nevertheless, in some cases the subjects, after having issued unreliable statements in front of police, come to the courtroom and confirm in a hearing the same testimony previously given. This is the less frequent situation because people who have lied during a hearing or in some other moment, have often the good sense of not repeating the crime twice and in front of the judge.

In fact, since these proceedings are aimed at verifying if the subject lied or not, they imply the development of investigative activities which make highly improbable the possibility of errors or bias in the identification of deceptive statements. Also because in the presence of any reasonable doubt about guilt, the defendant is acquitted. Furthermore, in the Italian Criminal Code an essential part of the crime is the so called ‘subjective element’, which refers to the fact that not only the not truthfulness of the statements has to be ascertained, but also the precise intent of the subject of deceiving the Judicial Authority. In the end, the outcome of these proceedings is a judgment that summarizes the facts and, when the defendant is found guilty, points out in a certain, organic and exhaustive way the lies which he told.

In this way it has been feasible to create DECOUR, a corpus of transcripts that contain the exact words pronounced by the subjects in the hearings, and about which it is possible to reliably know the truthfulness. In particular, in order to allow the study of deceptive language, DECOUR is made of hearings where the subjects have effectively been found guilty. To be more precise, in few cases the defendants have been acquitted, but merely for procedural and legal reasons: in every hearing which constitutes the corpus, there are lies told by the defendant, and these lies are recognized and clearly pointed out in the judgment.

## 3. Annotation

### 3.1. Mark up format

Hearings in Court are events strongly ritualized, with rules determined by the Code of Criminal Procedure. It means that the development of every hearing is highly regular, almost like in an experimental design, giving the opportunity of collecting data in relatively homogeneous conditions, even when the actors differ. The protagonist of each hearing is the subject who gives the testimony. He answers the questions posed by three other figures, who cannot be absent from any hearing: the judge, the public prosecutor and the defendant’s lawyer. Therefore, the considered testimonies have the form of a dialogue, in which at least four actors are present. It is possible that other actors intervene, for example more than one public prosecutor, or more than one defendant lawyer, or a lawyer for the victim of the crime, or a police officer: but these are less frequent cases.

Each text file that contains a testimony is transformed into XML format, with the aim of marking up actions and words of each participant. First, each XML has an **header** that

contains some meta-information about the testimony, such as place and date of the event, and about the speaker, such as his age, sex, place of birth and if known - unfortunately, not often - his level of instruction. The hearing properly said begins with an **introduction**: a formal part of the report which gives act to the introduction of the subject in front of the judge and, if needed, of his availability to answer the questions (to issue statements is an option for the defendant, but is a duty for the witnesses). Then, the real dialogue begins and each intervention of the different actors, delimited by the interventions of others participants, is marked as **turn**. Each turn can be constituted of one or more **utterances**, which are delimited by terminal punctuation marks: this is the atomic analysis unit of DECOUR. Into each turn, besides, some **action** carried out by the speaker can be inserted, according to what is minuted in the report. In the end, a **conclusion** can be present or not, with some last ordering of the judge or some ending formulas.

### 3.2. Coding scheme

Each utterance issued by the speaker receives a label, which concerns his degree of truthfulness. This annotation is carried out by hand, on the basis of the information found in the Court's final judgment. Obviously, between the white of the truth and the black of the falsity, there are wide gradations of gray, and the judgment, that describes the facts and points out the lies of the defendant, cannot give reason for each statement issued in the courtroom. This is the reason why the process of labeling the utterances of DECOUR followed a path which represented the research of a trade-off between opposite demands: the analytical representation of their degree of truthfulness and the achievement of a satisfying degree of agreement between different annotators, regarding this evaluation.

First step, DECOUR has been labeled according to the following annotation scheme.

**False** Utterance pointed out in the judgment as false, or of which the falsity is a logic consequence of some ascertained lie, are taken as false. For example, if the subject claims to have not been somewhere with someone, but actually he was, he also has to lie when he denies having known the same person.

Anyway, even though the judgment gives certain guidelines, it is not always easy to decide whether to assign to an utterance the label *false* or not. In fact, sometimes the meaning of linguistic behavior forces the focus on the function of the utterance, rather than on its literal sense. For example, from a theoretical point of view, questions do not represent any fact, therefore in strict sense they cannot be considered either true or false. But if a subject, pretending to not know a person, and asked "Do you know Mr. Rossi?", answers "Is Mr. Rossi the person in front of me?", the function of his answer/question is to generate in the judge a false representation of the reality, according to which he would not know Mr. Rossi. So, this utterance is considered a lie and labeled as false.

On the other end, some utterances which seem false, from a logical point of view are instead true. In one

proceeding, for example, a subject claimed to be an electronic engineer, while he had had a simple high school education. In front of the lawyer who was saying "You said you are an engineer...", the witness completed the sentence saying "electronic"<sup>2</sup>. Obviously, the fact that he was an engineer was false; but it was true that he had said that he was an engineer: then this answer was true, regarding the question posed.

**True** Utterances which are found coherent with the reconstruction of the facts contained in the judgment, are considered as true. Also the utterances which concern something not considered in the sentence, because they are uninfluential in respect to the investigated facts, are generally considered true.

For example, if the public prosecutor asks "For how long have you been married to Mrs. Bianchi?", and the subject answers "For eight years", this answer is considered true, even though the judgment says nothing about that, because there are no logical reasons to lie on this detail.

**Not reliable** Utterances which concern the investigated facts, but of which the truthfulness or deceptiveness is not established by the judgment, are considered not reliable. These utterances are related to some point about which the speaker could have some interest to lie, but the judgment does not provide the necessary information to evaluate them.

An interesting fact is that some judgments establish that the defendant was lying, when he had been claiming to not remember something. In these cases, the statements in which the subject says to not remember some specific event, are considered false. On the other hand, obviously, there are (many) proceedings in which it is not considered (or, at least, proved) as a lie the fact that the subject claims to not remember something. Also in this case, if the lack of memory is related to something that does not concern the topic of investigation, the utterance is considered sincere and labeled as true; otherwise, if it is related to the object of investigation, and the subject lying could defend his own interests, it is considered not reliable.

**True or not reliable** This class of utterances is similar to the not reliable ones. They are also related to the topic of investigation, and the judgment does not provide information about their truthfulness. Nevertheless, according to the event and/or on the basis of a weak connection with the interests that the speaker tries to defend, common sense induces one to believe that these utterances could be truth. The boundaries of this class of utterances resulted in being too subjective, and this caused problems of agreement between annotators, as will be discussed in the next subsection.

---

<sup>2</sup>In Italian, unlike the English language, often adjectives follow the noun (for example, "electronic engineer" is "ingegnere elettronico"). Therefore the lawyer was just waiting for the witness to complete his sentence.

**False or not reliable** This is the specular situation in respect to the previous point: the only difference is that the utterances seem to be false, even though their deceptiveness has not been clearly established by the inquiries.

**Undecidable** Utterances that, from a logical point of view, cannot be either true or false, are considered undecidable. Belonging to this class are questions, such as “Excuse me, can you repeat?”, but also of several utterances stopped in mid-sentence. This is also the case of utterances which have meta-communicative function, and regulate the relations between actors, such as “Now I’ll explain.” or “I would like to see you, if you were me...” and so on.

### 3.3. Agreement evaluation

The first studies carried out on DECOUR (Fornaciari and Poesio, 2011a,b) concerned preliminary analyses carried out on data collected in only three Courts, which represented the first nucleus of DECOUR. Since the study regarding the agreement between different annotators was not completed, these studies relied only on utterances held as surely true or false, having discarded the other ones.

The agreement study regarding the coding scheme described above was carried out employing three coders, each of whom marked 605 utterances, which meant about 20% of the final size of DECOUR. Kappa has been used as metric to evaluate their agreement (Artstein and Poesio, 2008), and its value was calculated under four different conditions, as follows:

**6 classes** The agreement was calculated on the previous coding scheme as it has been described;

**4 classes** The utterances marked up as *true or not reliable* were collapsed into the class *true*, while the *false or not reliable* utterances, in turn, were joined to the *false* ones. Then the four classes became *true*, *false*, *not reliable* and *undecidable*.

**3 classes** In this condition *true* and *true or not reliable* utterances, *false* and *false or not reliable*, and lastly *not reliable* and *undecidable* were respectively collapsed together into the classes *true*, *false*, *uncertain*.

**2 classes** In the last condition, the *false or not reliable* utterances were joined to the class *false*, while all the remaining utterances were collapsed into the generic class *not false*.

The values of Kappa under the different conditions are shown in Table 1. The values of K for two classes indicate a moderate to substantial agreement depending on whether we choose the interpretation of K values proposed by Carletta (1996) or that proposed by Landis and Koch (1977). Given that the fine-grained original annotation scheme was not suitable to reach a satisfying agreement between coders, in the end the whole DECOUR was annotated according to the only three collapsed classes: **true**, **uncertain** and **false**. Figure 1 shows an example of an XML file of DECOUR.

Table 1: Kappa values of the agreement studies.

| Classes evaluated | Kappa values |
|-------------------|--------------|
| 6 classes         | .40          |
| 4 classes         | .56          |
| 3 classes         | .57          |
| 2 classes         | .64          |

Table 2: Turns and utterances in DECOUR.

| Figure             | Turns | Utterances |
|--------------------|-------|------------|
| Speakers           | 2094  | 3015       |
| Public prosecutors | 1002  | 1323       |
| Judges             | 921   | 1201       |
| Defendant lawyers  | 388   | 527        |
| Police officers    | 3     | 4          |
| Tot.               | 4408  | 6070       |

Table 3: Labels of DECOUR’s utterances.

| Label       | Nr.  |
|-------------|------|
| True        | 1202 |
| Undecidable | 868  |
| False       | 945  |
| Total       | 3015 |

## 4. Corpus statistics

DECOUR has been collected in the Courts of four Italian towns: Bologna, Bolzano, Prato and Trento. It is constituted of 35 hearings, issued by 31 subjects. They appear 19 times as witnesses, 14 times as defendants, one time as expert witness and one time as victim of another crime. Their mean age at the time of the hearing is slightly higher than 36. 23 are men, 7 women and one transgender. The region of birth is northern Italy for 12 of them, center for 2, south for 9, while 8 subjects were foreigner but good Italian speakers. Lastly, the education is known only for six subjects: four of them having a high school education, one middle school and the last one elementary school.

Table 2 shows the number of turn and utterances of the participants in the hearings. While the utterances of other figures are not taken into consideration, the 3015 utterances of the speakers have been labeled as shown in Table 3: that is DECOUR contains 31.34% of false, 39.87% of true and 28.79% of uncertain utterances.

In terms of tokens, the size of DECOUR, with and without punctuation, is shown in Table 4. Punctuation marks are considered in blocks: this means, for example, that a single dot and the three dots of the ellipsis are both considered as a single token.

## 5. Conclusion

DECOUR is the first linguistic resource in Italian containing sincere and deceptive statements collected in a natural environment. Therefore, the only possible comparisons concerned previous studies related to different language. In

Figure 1: An example of XML file of DECOUR. Also lemmas and part-of-speech (POS) are embodied into the XML files.

```

<hearing>
  <header birtharea="N" birthplace="xxxxx" birthyear="xxxxx" court="BZ" day="xxxxx" idsub="xxxxx" month="xxxxx" name="xxxxx"
    nrdoxx="xxxxx" nrhear="xxxxx" sex="M" study="unk" typesub="defwit" typetest="false" year="xxxxx" yeardoss="03"/>
  <intro> giudice monocratico - dott.ssa xxxxx xxxxx viene introdotto il testimone ; questi viene avvertito dal giudice dei suoi obblighi e rende la
    dichiarazione ex art. 497 c.p.p. . fornisce le generalità : xxxxx xxxxx , nato a xxxxx il xx xxxxx xxxxx , ivi residente . </intro>
  <turn nrgen="1" nrpros="1" speaker="pros">
    <utterance class="x" nrgen="1" nrpros="1">
      lei nella primavera del 2001 ci può dire come ha conosciuto xxxxx xxxxx , in quali circostanze ?
    </utterance>
    <lemmas class="x" nrgen="1" nrpros="1">
      lei nella primavera del @card@ ci potere dire come avere conoscere <unknown> <unknown> , in quale circostanza ?
    </lemmas>
    <pos class="x" nrgen="1" nrpros="1">
      PRO:pers ARTPRE NOUN ARTPRE NUM CLI VER2:fin VER:infi WH AUX:fin VER:ppast NOUN ADJ PUN PRE DET:wh NOUN SENT
    </pos>
  </turn>
  <turn nrgen="2" nrsub="1" speaker="defwit">
    <utterance class="uncertain" nrgen="2" nrsub="1">
      adesso non mi ricordo come l' ho conosciuto , comunque ci siamo conosciuti ...
    </utterance>
    <lemmas class="uncertain" nrgen="2">
      adesso non mi ricordare come l' avere conoscere , comunque ci essere conoscere ...
    </lemmas>
    <pos class="uncertain" nrgen="2">
      ADV NEG CLI VER:fin WH CLI AUX:fin VER:ppast PUN WH CLI AUX:fin VER:ppast SENT
    </pos>
  <turn nrgen="3" nrsub="2">
    <utterance class="uncertain" nrgen="3" nrsub="2">
      non mi ricordo , in giro , al xxxxx anche , perché prendevo il metadone tempo fa .
    </utterance>
    <lemmas class="uncertain" nrgen="3">
      non mi ricordare , in giro , al <unknown> anche , perché prendere il metadone tempo fa .
    </lemmas>
    <pos class="uncertain" nrgen="3">
      NEG CLI VER:fin PUN PRE NOUN PUN ARTPRE NOUN ADV PUN WH VER:fin ART NOUN NOUN ADV SENT
    </pos>
  <turn nrgen="4" nrsub="3">
    <utterance class="uncertain" nrgen="4" nrsub="3">
      adesso sono due anni che sono a posto , quasi due anni .
    </utterance>
    <lemmas class="uncertain" nrgen="4">
      adesso essere due anno che essere a posto , quasi due anno .
    </lemmas>
    <pos class="uncertain" nrgen="4">
      ADV VER:fin DET:num NOUN CHE VER:fin PRE NOUN PUN ADV DET:num NOUN SENT
    </pos>
  </turn>

```

Table 4: DECOUR’s size.

| Utterances | Tokens      |       |                |       |
|------------|-------------|-------|----------------|-------|
|            | With punct. |       | Without punct. |       |
|            | Mean        | Tot.  | Mean           | Tot.  |
| True       | 12.86       | 15456 | 10.67          | 12847 |
| Uncertain  | 12.02       | 10439 | 9.99           | 8669  |
| False      | 16.85       | 15924 | 14.15          | 13376 |
| Total      |             | 41819 |                | 34892 |

fact, the first preliminary analyses (Fornaciari and Poesio, 2011a) were aimed at replicating a well known study of Newman et al. (2003), in which lexical features were employed in order to classify texts as truthful or untruthful. The results showed that it was possible to train models with the ability to classify true and false utterances with a degree

of accuracy at 72/73%, from a baseline of 60%. A second study, in which surface features were employed with the same goal, analogous results were found (Fornaciari and Poesio, 2011b). A third study (Fornaciari and Poesio, 2012) will show that removing outliers from the corpus improves the models’ performance. This last paper, currently in press, is the first in which the analyses are carried out not on the subset of utterances held as true or false, but on all the utterances of DECOUR, divided into the two classes *false* and *not false*. So grouped, the data are evidently more noisy, and the models’ performance suffers from this. Nevertheless, the overall accuracy is 6% above the baseline.

Therefore the first results of the analyses on DECOUR suggest that, even though the analysis units are extremely short texts, stylometric techniques can be successfully applied in order to train models aimed to distinguish deceptive from truthful language.

## Acknowledgements

To create DECOUR has been very complex, and it would not have been possible without the kind collaboration of a lot of people. Many thanks to Dr. Francesco Scutellari, President of the Court of Bologna, to Dr. Heinrich Zanon, President of the Court of Bolzano, to Dr. Francesco Antonio Genovese, President of the Court of Prato and to Dr. Sabino Giarrusso, President of the Court of Trento.

## References

- Artstein, R. and Poesio, M. (2008). Inter-coder agreement for computational linguistics. *Comput. Linguist.*, 34(4):555–596.
- Bachenko, J., Fitzpatrick, E., and Schonwetter, M. (2008). Verification and implementation of language-based deception indicators in civil and criminal narratives. In *Proceedings of the 22nd International Conference on Computational Linguistics - Volume 1, COLING '08*, pages 41–48, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Bond, C. F. and De Paulo, B. M. (2006). Accuracy of Deception Judgments. *Personality and Social Psychology Review*, 10(3):214–234.
- Carletta, J. (1996). Assessing agreement on classification tasks: the kappa statistic. *Comput. Linguist.*, 22(2):249–254.
- Fitzpatrick, E. and Bachenko, J. (2009). Building a forensic corpus to test language-based indicators of deception. *Language and Computers*, 71(1):183–196.
- Fornaciari, T. and Poesio, M. (2011a). Lexical vs. surface features in deceptive language analysis. In *Proceedings of the ICAIL 2011 Workshop Applying Human Language Technology to the Law, AHLTL 2011*, pages 2–8, Pittsburgh, USA.
- Fornaciari, T. and Poesio, M. (2011b). Sincere and deceptive statements in italian criminal proceedings. In *Proceedings of the International Association of Forensic Linguists Tenth Biennial Conference, IAFL 2011*, Cardiff, Wales, UK.
- Fornaciari, T. and Poesio, M. (2012). On the use of homogenous sets of subjects in deceptive language analysis. In *Proceedings of the EACL 2012 Workshop on Computational Approaches to Deception Detection, EACL 2012*. In press.
- Hancock, J. T., Curry, L. E., Goorha, S., and Woodworth, M. (2008). On Lying and Being Lied To: A Linguistic Analysis of Deception in Computer-Mediated Communication. *Discourse Processes*, 45(1):1–23.
- Landis, R. J. and Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174.
- Levine, T. R., Feeley, T. H., McCornack, S. A., Hughes, M., and Harms, C. M. (2005). Testing the Effects of Nonverbal Behavior Training on Accuracy in Deception Detection with the Inclusion of a Bogus Training Control Group. *Western Journal of Communication*, 69(3):203–217.
- Newman, M. L., Pennebaker, J. W., Berry, D. S., and Richards, J. M. (2003). Lying Words: Predicting Deception From Linguistic Styles. *Personality and Social Psychology Bulletin*, 29(5):665–675.
- Strapparava, C. and Mihalcea, R. (2009). The Lie Detector: Explorations in the Automatic Recognition of Deceptive Language. In *Proceeding ACLShort '09 - Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*.
- Zhou, L., Burgoon, J. K., Nunamaker, J. F., and Twitchell, D. (2004). Automating Linguistics-Based Cues for Detecting Deception in Text-based Asynchronous Computer-Mediated Communication. *Group Decision and Negotiation*, 13(1):81–106.
- Zhou, L., Shi, Y., and Zhang, D. (2008). A Statistical Language Modeling Approach to Online Deception Detection. *IEEE Transactions on Knowledge and Data Engineering*, 20(8):1077–1081.