

Detecting deception in conversations

Ioana-Raluca Zaman

Politehnica University of Bucharest
313 Splaiul Independentei,
Bucharest, Romania
ioana.raluca.zaman@gmail.com

Ștefan Trăușan-Matu

Politehnica University of Bucharest
313 Splaiul Independentei,
Bucharest, Romania
and
Research Institute for Artificial
Intelligence
and
Academy of Romanian Scientists
stefan.trausan@cs.pub.ro

ABSTRACT

Interaction with other people is part of anyone's life, and implicitly comes with the risk of deception. On average, a person is lied up to several dozen times a day, the main purpose being manipulation. Deception is defined as a psychological act of misleading or trying to persuade someone to believe certain facts. One manner to understand deception is *via* emotions and sentiments of the participants in conversation. This paper's aim is to combine the information extracted from the text with the help of techniques specific to Natural Language Processing with notions of psychology. Even if lies are not a mandatory element, in most cases they play a crucial role in manipulation. Therefore, this project classifies lies based on the emotions extracted from a message, but also from the replies to it and using two alternatives of the Naïve Bayes algorithm. Moreover, the project will focus both on actual lies and suspected lies. The corpus utilized for analyze is *Deception in Diplomacy Dataset* from Cornell University ConvoKit.

Author Keywords

natural language processing; deception; chat; sentiments; emotions

ACM Classification Keywords

I.2.7 Natural Language Processing: Text analysis.

DOI: 10.37789/rochi.2021.1.1.5

INTRODUCTION

Everyone interacts with other people, now in an increasing degree mediated by technology, the connections between people being formed through interactions. The basis of a real connection is trust, but it can be easily destroyed if signs of deception appear. This is a complex phenomenon which can have various intentions (e.g., a politician trying to convince the public that he deserves to win the election or a child trying to convince his parents that he did not eat the cake) and therefore it is quite difficult to detect. Besides the well-known Polygraph (i.e., device that detects lies based on cardiovascular, electrodermal or respiratory changes), there are other techniques specialized in this such as: studying

brain's activity using electrodes, facial expressions, body language, voice changes or linguistic characteristics [5]. Nowadays, communication through online messages is probably the most used (e.g., in 2020 approximately 2.77 billion people with a mobile phone used an application for messaging [8]), hence deception occurs more and more often in these circumstances and because of this, it should be studied in such an environment (i.e., online dialogue). Deception detection is a problem analyzed from several points of view such as polygraph, facial expressions or body language. Study of deception from the perspective of NLP is a challenge, having access only to the textual cues.

A dialogue is a dynamic act in which at least two people participate, exchange ideas and influence each other. Therefore, analyzing the emotions and sentiments of the participants in their conversation could improve the detection of deception.

Even if a person may feel a multitude of emotions at a time, the focus of the project will be on the eight basic emotions proposed by Robert Plutchik [10]. The emotions are illustrated in a wheel-like diagram in which every primary emotion has an opposite, these being: *joy-sadness*, *trust-disgust*, *fear-anger* and *surprise-anticipation*. Naturally, the emotions experienced by a manipulator are quite the opposite of those experienced by the one being deceived (e.g., *disgust* or *sadness* respectively *trust* or *joy*).

From the point of view of NLP, the emotions can be quantified in several ways, one of them employing the help of lexicons for Sentiment Analysis. Lexicons provide a score for each emotion or sentiment which is a more accurate method than others (i.e., detecting emotions through punctuation marks).

Another aspect the relationship between the manipulator and the manipulated because deception is an act involving two or more people. In most cases, there is a lack of balance between individuals in hierarchical social positions (e.g., *politician/citizen* or *parent/child*), knowledge (e.g., *expert/non-expert*) or power/domination (e.g., in a game *player who leads/player who is led*) [4]. For the corpus used in this project, this aspect can be studied *via* the difference between players' scores.

STATE OF THE ART

Deception cues

When people make deceptive statements, they reference themselves less, often using the third person, trying to disassociate, and having a distancing language (e.g., *that man*). Because they are always thinking about what to say next to achieve their goal, manipulators tend to use simple terms, but long phrases, sometimes trying to insert unnecessary and irrelevant words to distract the attention and sometimes have pauses in speech (e.g., ‘*I want to study and get an A on the exam.*’ vs ‘*I want to get organized and make a study plan with the help of which I can prepare for the mathematics exam I will have on Wednesday to get the highest grade.*’). Moreover, people who lie tend to be more hesitant, use more qualifiers (e.g., *maybe*) and try to avoid really answering questions or because they know they are lying use words suggesting uncertainty (e.g., *like, almost, kind of*). Hatzivassiloglou and McKeown [3] focused on the role of conjunctions in classification of adjectives and acquired a precision over 90%. Their central idea is that adjective conjoined by *and* or *even* or *are* often similar, while those conjoined by *but* have different orientations, such as “*beautiful and smart*” and “*smart but arrogant*”.

Lexicons for Sentiment Analysis

A lexicon for Sentiment Analysis can be defined as a list of words where the words are grouped into categories. NRC Word-Emotion Lexicon known as EmoLex is one of them and it is formed based on the Plutchik [9] diagram and two sentiments (*positive* and *negative*). Another tool for sentiment analysis is Valence Aware Dictionary and sEntiment Reasoner known as VADER, this lexicon produces four metrics: *positive, negative, neutral, and compound score*. In addition to emotions and sentiments, important features are also *subjectivity* and *polarity*, which can be discovered with TextBlob package. Subjectivity refers to personal expressions such as opinions, feelings or beliefs. Polarity measures the presence of emotions (e.g., *positive, or negative*) in a phrase.

Related work

The work of Danescu-Niculescu-Mizil et al. [2] offers state-of-the-art methods using logistic regression and neural models. Using the same corpus, the authors emphasize both what is *intended* as a lie and what is *perceived* as a lie, the research has tended to focus on tactics from Cialdini and Goldstein’s [1] paper (e.g., *reciprocity*) and internal attitudes (e.g., *sarcasm*). In both approaches, the result has four possibilities: Deceived (lie perceived as truth), Caught (lie perceived as lie), Straightforward (truth perceived as truth) or Cassandra (truth perceived as lie). In the first approach, researchers used betrayal specific linguistic Harbingers from the paper of Nicolae et al. [6] for logistic regression model. Having better results, the neural attempt takes into consideration both the content and previous context by using a LSTM (Long Short-Term Memory). Moreover, BERT (Bidirectional Encoder Representations from Transformers)

embeddings were utilized, but no remarkable progress was observed. To sum up, the work of Danescu-Niculescu-Mizil et al. has notable results, the best macro *F1* score for actual and suspected lies being 57.2 respectively 54.3.

Zhou and co-workers [7] studied deception in TA-CMC (Text-based Asynchronous Computer Mediated Communication) taking into consideration 27 linguistic features grouped in the following categories: affect, complexity, diversity, expressivity, informality, non-immediacy, quantity, uncertainty and specificity. Their study demonstrates that lies are longer than truths and more expressive and informal. Moreover, the deceivers’ messages were less diverse, complex in punctuation and self-referenced.

Corpus

The dataset utilized in this project is suitable for studying deception since it illustrates long-lasting dialogues between players in the context of a political negotiation game, thus replicating a real situation in which lies and attempts at manipulation are among many true statements. The strategic board game can be played by two to seven players, each of them representing one of the top seven European powers of the 20th century (Austria-Hungary, France, Germany, Great Britain, Italy, Turkey, and Russia). Beginning in 1901, each year of the game is separated in two seasons: Spring and Fall and each season has two phases: Negotiation and Movement. Negotiation is the phase when manipulation can be studied, when the players form alliances or disinform the enemies, following by the Movement phase where each player writes the actual orders and these are simultaneously performed. In the Diplomacy game, every country has two units: an army and a fleet and all of them have the same strength and in a province (which can be water, coastal or inland) can be only one unit. On the map, there are 34 supply centers with which armies for winning the game can be built and the score of a player is represented by the number of supply centers which (s)he owns. Over 17,000 messages from conversations on Discord are annotated by sender and receiver. Online conversations are done by a diverse group of players consisting both of leading and experienced players and inexperienced, but dedicated players. They communicate indirectly *via* a bot which forwards the message only after the sender annotates the message. Each player has a different strategy but there are some that are frequently used, such as keep communication with everyone (enemies or friends), do not be extremely successful from the beginning because that will make you the target, offer help to players who cannot attack you or look for allies on the opposite side of your enemy [9].

IMPLEMENTATION DETAILS

Differences between truths and lies

By analyzing the messages from conversations, the average number of words in both true utterances and lies confirms the

fact that manipulators tend to utilize longer phrases by trying to insert needless and irrelevant words to distract attention.

Average length	Truth	Lie
Initial message	20.35	30
Normalized message	8.91	12.8

Table 1. Statistics on average number of words in messages from train data

The two distributions of emotions in lies and sincere messages are quite similar, but there still are some differences. For example, in lies *surprise* and *joy* have higher scores, whereas in truth messages *positive* and *fear* are higher. In addition to this comparison, some differences can be noticed between the scores of the replies, especially if the players suspect or not if they are lied. In the suspected lies chart, fear has a higher score, possibly due to uncertainty. Contrariwise, trust has a greater score, this being explained by the fact that most lies are believed. On the other hand, receivers who do not suspect a lie are more *surprised* and *joyful*, but with a higher *anticipation* score. One explanation consists in the fact that the manipulators are trying to convince the receivers to accept alliances that could help them win.

The difference between scores is representative for the power ratio between players. People tend to manipulate more when they are in a dominant position because they consider themselves in control. As Figure 1 shows, most lies are when the player is almost in control (i.e., the difference between scores is bigger than -1). Also, it can be seen that the greater the difference is, the lower the number of lies is. One explanation for this fact is that people no longer lie because they do not feel threatened anymore.




Figure 1. Crosstab representation of all lies in train dataset according to difference between players' scores

Another aspect is the actual player's score. Figure 2 illustrates that most lies are when the score is in the range [3, 10], highlighting that players did not lie when they felt led or that they had nothing to gain anymore.

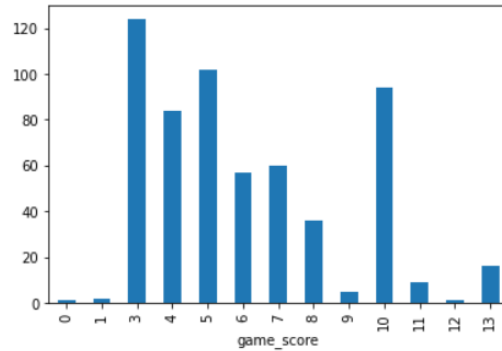


Figure 2. Crosstab representation of all lies in train dataset according to players' scores

Another aspect about the emotions from the dataset is the correlation between them. Correlation coefficients measure the association between two features from a dataset and have values in the range [-1, 1]. Correlation coefficients in this case are approximately only in the range [-0.2, 1]. As highlighted in Figure 3, there is a moderate positive correlation between *negative*, *sadness* and *disgust* with *fear* and *anger* and between *positive*, *trust* and *surprise* with *joy*. These correlations can be explained by the fact that there are words that are labeled by multiple emotions.




Figure 3. Correlation matrix between emotions from EmoLex in messages illustrated as a heatmap

In addition to studying the differences between truths and lies, the project also classifies them. The source code of the program has been written in Python and has four main sections. Each section will be presented in detail in the following paragraphs.

Text preprocessing

In this part each message is adjusted to make scores of emotions and sentiments extracted from them as relevant as possible. Initially, in the JSON each object represents a game dialog. Hence, for the individual processing of each message, the dataframe resulting from the JSON file has been converted in a new one in which each row corresponds to a message. All operations on dataframes were made using

pandas library. For example, here is a reply from the train dataframe in a JSON format.

```
{
  "speakers": "italy",
  "message": "Germany!\n\nJust the person...",
  "sender_label": true,
  "receiver_label": true,
  "game_score": 0,
  "game_score_delta": 3,
  "absolute_message_index": 74,
  "game_id": 1,
  "year": 1901
}
```

Figure 4. Example of a reply from the game containing the message and its metadata

```
{
  "message": "Germany!\n\nJust the person...",
  "sender_label": true,
  "receiver_label": true,
  "game_score": 0,
}
```

Figure 5. Example of a reply from the game after the unused fields were removed

After the procedures applied on dataframe, each message was tokenized. The procedure of tokenization is the one which segments the textual input into tokens. For instance, the first part of the message from the example above becomes ‘Germany!\n\nJust the person I want to speak with. I have a somewhat crazy idea that I’ve always wanted to try with I/G, but I’ve never actually convinced the other guy to try it.’ → [‘Germany’, ‘!’, ‘\n’, ‘\n’, ‘Just’, ‘the’, ‘person’, ‘I’, ‘want’, ‘to’, ‘speak’, ‘with’, ‘.’, ‘I’, ‘have’, ‘a’, ‘somewhat’, ‘crazy’, ‘idea’, ‘that’, ‘I’, ‘’, ‘ve’, ‘always’, ‘wanted’, ‘to’, ‘try’, ‘with’, ‘I’, ‘/’, ‘G’, ‘,’, ‘but’, ‘I’, ‘’, ‘ve’, ‘never’, ‘actually’, ‘convinced’, ‘the’, ‘other’, ‘guy’, ‘to’, ‘try’, ‘it’, ‘.’ J). After that, the following modification have been made, performing the normalization process:

- Every non-alphabetic character was replaced by “spaces” using RegEx package.
- Every letter was converted to lowercase.
- Stop words from English vocabulary are eliminated.
- The names of the countries in the game are eliminated because they were very frequently used and did not have helpful information.

Stop words (e.g., *the*, *of*, *to*) are a list of frequently utilized words, but without having valuable information that are filtered out from corpus. Built on mathematical statistics, Zipf’s empirical law claims that the relation between the frequency of a word appearing in a corpus and its rank of occurrence is inverse proportionality. Removing stop word will result in a more representative text, a faster model and additionally can increase the performance.

After the normalization, lemmatization techniques were applied using the nltk package. In lemmatization, a word is replaced by the root word or a word with similar context by extracting the lemma from it (e.g., ‘*oxen*’ → ‘*ox*’, ‘*best*’ → ‘*good*’. After all these steps, the phrase from the example is transformed into: ‘*person want speak somewhat crazy idea always wanted try g never actually convinced guy try*’.

Obtaining the replies

One of the key ideas of this project is that the algorithm does not only consider the scores of emotions for messages, but also those of responses to them. Because the dataset consists of online written dialogues, each message can be answered with one or more replies. In case that a message has been answered with several replies, they were concatenated in a single response. Messages and their replies were saved in a dictionary in which the key was represented by message and the value by the replies concatenated. For the message analyzed before, its normalized reply is: ‘*whet appetite suggestion*’, the original message being ‘*You’ve whet my appetite, Italy. What’s the suggestion?*’.

Computation of emotions’ scores

Emotions and sentiments are the points of interest in this approach. Scores of emotions and sentiments were computed using Sentiment Lexicons.

First lexicon used is NRC Word-Emotion Lexicon known as NRCLex (<https://pypi.org/project/NRCLex/>), being the one with the most words and providing the most features for the model, more exactly 20. Using this, scores for eight emotions (*anger*, *disgust*, *anticipation*, *joy*, *fear*, *trust*, *sadness* and *surprise*) and two sentiments (*negative* and *positive*) were computed based on their frequencies in text. In Table 2 are presented the scores for the above example

Anger	Anticipation	Disgust	Fear	Joy
0.12	0.12	0	0.12	0.08
Sadness	Surprise	Trust	Positive	Negative
0.12	0.04	0.12	0.12	0.16

Table 2. Scores of emotions and sentiments computed using NRCLex

Other significant features were extracted using TextBlob package, these being *polarity* and *subjectivity* for both messages and their replies. With a few exceptions, lies have both greater *polarity* and *subjectivity* than true statements. For example, for the message polarity and subjectivity are -0.08 and 0.59, respectively.

Last lexicon used is VADER from which *positive*, *negative*, *neutral*, and *compound score* were computed for messages and their replies. Liars are more *positive* and a little less *neutral*. The replies given by players who suspected a lie are more *negative* and *neutral*. Using the same example, the scores are: positive – 0.23, negative – 0.262, neutral – 0.508 and compound score – 0.193.

Classification

The purpose of this project is to classify the messages with the help of emotions scores. As it is one of the most popular classification algorithms in NLP problems, the following two models of the algorithm Naïve Bayes were used: Gaussian Naïve Bayes and Complement Naïve Bayes. The approach is centered on detecting two types of messages – actual lies (message is labeled as a lie by sender) and suspected lies (message is labeled as a lie by receiver) with two types of Naïve Bayes classifiers.

Before the classification, three methods were applied to dataset to optimize the algorithms:

- Data preprocessing
- Scaling features
- Feature selection
- Hyperparameter tuning

One preprocessing technique is scaling features to a range (e.g., [0, 1], [-1, 1]), producing robustness and preserving zero entries in sparse data. For the Gaussian Naïve Bayes approach, the `MaxAbsScaler` from `sklearn.preprocessing` package was applied and this scaler operates via dividing every value by the largest maximum value in each feature, converting the data in range [-1, 1]. Due to the fact that Complement Naïve Bayes does not accept negative values, `MinMaxScaler` also from `sklearn.preprocessing` package was utilized.

The second procedure for classification optimization is feature selection which involves choosing just the relevant features. This selection prevents overfitting and improves both accuracy and classification time. For selection, `SelectKBest` method from `sklearn.feature_selection` package was utilized having as score function parameter `f_classif`.

Lastly, hyperparameter tuning was made using `GridSearchCV` from `sklearn.model_selection`. `GridSearchCV` tests all values given to the parameters and selects the ones that have best results. By the reason that both `GaussianNB` and `ComplementNB` have a modest number of hyperparameters, the only optimization found was for alpha smoothing parameter from `ComplementNB`.

RESULTS

Considering the imbalanced binary dataset, the metrics will use a macro average and the representative metric will be F1 score. The values of macro F1 for Actual Lies task has as its baseline Random and Majority Class (i.e., the model always chooses the class that is in majority in train set – in this case Truth class) approaches. Naturally, the Major Class approach achieves score 0 for Lie F1 and for Suspected Lies task, the Human score is no longer considered. Human score is computed using the receivers' labels from corpus.

For detecting both types of lies using the two algorithms, the following alternatives were tested: Emotions, Emotions Delta Scores, Emotions + Scores, Emotions + Delta Scores + Scores.

Since `ComplementNB` classifies according to the probability of an element not belonging to a class, it is understandable why this approach has better results on F1 Lie than the Gaussian approach while Gaussian approach has better results for macro F1.

An observation about the results is that the addition of Scores or Delta Scores did not improve the scores significantly, excepting the Gaussian approach for Actual Lies.

A notable aspect is which type of lies were classified correctly by the two Naïve Bayes models and by human. The lies are divided into four categories: detected by both players and models, detected only by models, detected only by players, undetected by none.

The lies detected both by players and models have a general positive and flattering attitude and are quite long, for the Gaussian approach the average number of words in a message by both being 35 and for the Complement approach being 36.6 (e.g., *'You're funny. Why is giving up Sev "way too dangerous? It's not like me having it gives me access to very much of your territory... it is Russian territory after all. You are already a very powerful country, and if you do as I suggest, moving into Gre and Serbia, you could have nine supply centers. If you give me Sev, I will support you into Bud'*).

Regarding the messages classified correctly only by models, in most cases they are appropriately excuses (e.g., *'Sorry for the stab earlier.'*, *'Oh sorry. I've been real busy'*) and plans for the future (e.g., *'Let's see...'*, *'excellent, it's a plan!'*).

Messages detected only by players, models encounter difficulties in identifying lies that can be caught from the other discussions. In most cases, the lies involve other countries (e.g., *'Germany (who says is allied with Austria) is telling me...'*, *'Hey what's the deal with France? They're talking about an alliance between us three to counter Russia and Italy?'*) and are about their strategies.

Players and all models were deceived by the following category of messages: the concise, simple and some of them really short (e.g., *'That works for me'*, *'I'm with Austria. Going for Turkey.'*). Moreover, the messages which contain a personal involvement (e.g., *'I was hoping...'*, *'I am very worried...'*) or the word *'trust'* or *'honest'* are classified as truths.

Actual Lies							
Method \ Algorithm	Random	Majority Class	Emotions	Emotions + Delta Score	Emotions + Scores	Emotions + Scores +	Human
Macro F1							
GaussianNB	39.8	47.8	49.46	53	53.45	52.7	58.1
ComplementNB			44.76	44.57	45.16	44.83	
F1 Lie							
GaussianNB	14.9	0	5.26	13.52	15.81	16.2	22.5
ComplementNB			18.37	18.56	18.38	18.34	

Table 3. Test results for both Naïve Bayes algorithms for Actual lies

Suspected Lies						
Method \ Algorithm	Random	Majority Class	Emotions	Emotions + Delta Score	Emotions + Scores	Emotions + Scores +
Macro F1						
GaussianNB	38.3	48.3	51.16	51.75	51.7	51.52
ComplementNB			42.7	42.8	43.32	43.32
F1 Lie						
GaussianNB	11.8	0	10.8	10.8	10.55	10.24
ComplementNB			11.25	11.36	11.2	11.14

Table 4. Test results for both Naïve Bayes algorithms for Suspected lies

The central paper for the state of the art is ‘It Takes Two to Lie: One to Lie, and One to Listen’ of Danescu-Niculescu-Mizil et al. [2], which is based on the same dataset from Diplomacy game. In the next horizontal bar charts will be shown a comparison between the test results of the research presented herein and the ones from the paper of Danescu-Niculescu-Mizil et al. The following figures illustrates the results for Macro F1 and Lie F1, classifying Actual and Suspected lies.

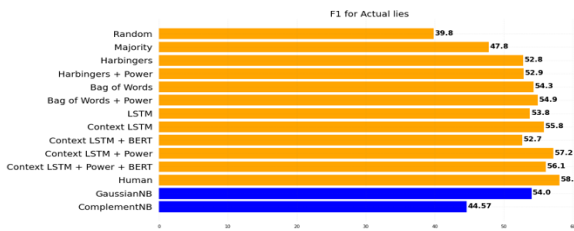


Figure 6. Test results for both Naïve Bayes algorithms compared to results from ‘It Takes Two to Lie: One to Lie, and One to Listen’ for Macro F1 score for Actual Lies task.




Figure 7. Test results for both Naïve Bayes algorithms compared to results from ‘It Takes Two to Lie: One to Lie, and One to Listen’ for Lie F1 score for Actual Lies task.




Figure 8. Test results for both Naïve Bayes algorithms compared to results from ‘It Takes Two to Lie: One to Lie, and One to Listen’ for Macro F1 score for Suspected Lies task.

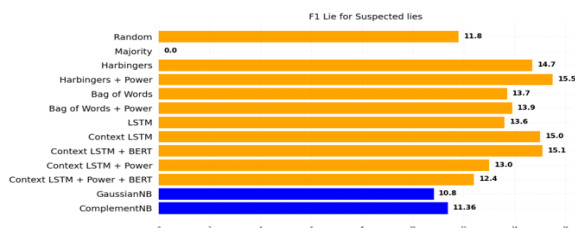


Figure 9. Test results for both Naïve Bayes algorithms compared to results from ‘It Takes Two to Lie: One to Lie, and One to Listen’ for Lie F1 score for Suspected Lies task.

CONCLUSIONS AND FURTHER WORK

The suggested solution aims to detect deception in conversations by analyzing the emotions for both the speaker and the receiver. Because of the imbalanced dataset, F1 score is the principal metric used, more precisely Macro F1 and Lie F1.

The results can be improved by making the model more general (i.e., to detect lies from several categories). The modifications that could be incorporated in the proposed solution are exploiting more cues of the attitude of the players and text properties, taking into account the context or using more suitable algorithms. For example, the rank that a player occupies at a given time could be a cue. Furthermore, there are indicators of certain emotions that can be extracted by methods other than lexicons. For instance, pauses in speech can indicate a lie by a thinking or even guilty attitude.

So far, the emphasis has been on detecting deception in singular elements (e.g., news, online posts) that are not depending on the context and can be more easily studied. Manipulation is often performed in a direct human interaction (i.e., a dialogue), therefore my solution is useful and can be utilized in cases such as: chat applications to ensure security, analysis of negotiation discussions or even as an interrogation tool. Moreover, the results can be used for analyzing disinformation in online environment such as forum discussions or fake reviews or posts.

REFERENCES

1. Cialdini, R. B., Goldstein, N.J., 2004. Social influence: Compliance and conformity, *Annual Review of Psychology*, 55:591–621
2. Danescu-Niculescu-Mizil, C., Peskov, D., Cheng, B., Elgohary, A., Barrow J., Boyd-Graber J., (2020) It Takes Two to Lie: One to Lie, and One to Listen, *58th Annual Meeting of the Association for Computational Linguistics*, pages 3811–3854
3. Hatzivassiloglou, V. and McKeown, K. (1997). Predicting the semantic orientation of adjectives. *35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics*
4. Maillat, D., & Oswald, S. (2009). Defining Manipulative Discourse: The Pragmatics of Cognitive Illusions, *International Review of Pragmatics*, 1(2), 348-370.
5. National Research Council (2003), *The Polygraph and Lie Detection*. Washington, DC: *The National Academies Press*.
6. Niculae, V., Kumar, S., Boyd-Graber, J., Danescu-Niculescu-Mizil, C., 2015. Linguistic harbingers of betrayal: A case study on an online strategy game in *Proceedings of the Association for Computational Linguistics*.
7. Zhou, L., Burgoon, J., Nunamaker, J., & Twitchell, D.P. (2004). Automating Linguistics-Based Cues for Detecting Deception in Text-Based Asynchronous Computer-Mediated Communications. *Group Decision and Negotiation*, 13, 81-106.
8. <https://www.statista.com/statistics/483255/number-of-mobile-messaging-users-worldwide/#statisticContainer>, last accessed 28 August 2021
9. <https://www.ultraboardgames.com/diplomacy/tips.php>, last accessed 15 April 2021
10. <https://www.6seconds.org/2020/08/11/plutchik-wheel-emotions/>, last accessed 10 July 2021