

Enunciation and topic/comment structure: the offensive replies to Pope Francis' tweets¹

Francesco Galofaro

Dipartimento di filosofia e scienze dell'educazione, Università di Torino
francesco.galofaro@unito.it

Zeno Toffano

Laboratoire des signaux et systèmes, CentraleSupélec, Université Paris-Saclay
zeno.toffano@centralesupelec.fr

Abstract Sentiment analysis is an automatised technique of analysis aimed to measure the “polarity” and the “subjectivity” of large corpora of messages. The case study of the present paper consists of a selection of Pope Francis' tweets on ecological, social, religious themes and the relative polemic replies. In the degree of agreement/disagreement in response to a tweet, the referential function is not relevant; the emotive and conative functions prevail. The political strategies aimed at corroborating or refuting claims in terms of “fact checking” seem not relevant to these forms of communication based on personal enunciation, on the relation between the two simulacra “me” and “you”, and on the manifestation of one's own comment with respect to a topic. Furthermore, the techniques aimed at detecting the presence of hate speeches to apply, possibly, a precautionary censorship are lexical-sensitive, and fail to consider the context in which words co-occur. Finally, the paper presents a technique of analysis based on quantum information retrieval which can provide new insights on the relation between hashtag, address sign, topic, and reply.

Keywords: semiotics, quantum semantics, information retrieval, machine learning, sentiment analysis, hate speech, conspiracy theory.

Received 03/02/2022; accepted 23/04/2022.

0. Introduction

The analysis of tweets is an important benchmark for semiotic to test its views on meaning and to cooperate with other social sciences. «The theme of using statements

¹This paper is part of the research project NeMoSanctI (New Models of Sanctity in Italy (1960s–2000s)—A Semiotic Analysis of Norms, Causes of Saints, Hagiography, and Narratives), which has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 757314).

and comments published on Twitter (...), in fact, can represent an interesting starting point to reflect on semiotics and on how to conduct researches» (Santangelo 2015: 67). However, there are not many semiotic studies proposing a critical point of view on machine learning-based techniques of classification, which are indeed recent. Seven years ago, the major Italian semiotic expert in the field of semiotics and communication consultancy wrote:

Sentiment analysis (...) concretely consists in the automatic detection of potentially positive/negative statements, which express an agreement or disagreement in relation to the investment of the brand in its proposal. However, none of the algorithms in place can automatically evaluate the positive/negative belonging of a given word, which is evidently given by the contest. The recourse is invariably to qualitative analysis which, however, gives up in the face of numerous posts which can be, for the most active brands, in the order of a few thousand. Thus, paradoxically, we like sentiment analysis because it represents the last barrier to an interpretative practice that is intended to be fully automated (Giulia Ceriani 2015: 48, translated by the author)

As the present paper would like to prove, Ceriani's judgment on the algorithms used to classify thousands of tweets is still accurate nowadays, and not only as regards sentiment analysis, but also other sensible fields such as fact checking and hate speech detection: the training techniques are still very sensitive to the presence of lexical units, independently from context in which they occur. Nevertheless, only seven years after Ceriani's paper, these techniques have become practically ubiquitous, in spite of vague definitions of notions such as "subjectivity", "hate speech", "truth" – see section 2. In this framework, the paper aims to a better understanding of the relation between tweet, hashtag, and offensive replies. It will focus on polemical, offensive replies, investigating their relations to the morphological markers # and @, which manifest at the morphosyntactic level the topic and the sender/receiver of the message. In order to analyse the context of the lexical choices, the paper will use a technique developed by the authors, merging semiotics and quantum information retrieval.

1. Corpus and machine learning classification

The considered corpus consists of 2164 English comments in reply to Pope Francis's tweets on ecological and social themes, retrieved on Twitter and referring to the years 2015-2021. While the interesting features of the Pope's communication have been the object of many semiotic analysis (Peverini and Lorusso 2017), the present study will focus on the reactions of the readers, and, in particular, on their *polarisation*. Traditional machine learning-based classification algorithms will be evaluated and criticised from a semiotic point of view, and compared to different techniques of semantic analysis inspired by semiotics and based on quantum information retrieval (Van Rijsbergen, C.J. 2004; Melucci 2015).

Two Python libraries trained on the basis of machine learning techniques, *TextBlob* and *Hatesonar*, have been used to label the corpus to identify polemical answers to the Pope. In particular, *TextBlob*² performs the sentiment analysis of each message in terms of *polarity* and *subjectivity*. According to the documentation of the library, the polarity score is a float within the range (-1.0, + 1.0), while subjectivity is a float within the range (0.0, 1.0) where 0.0 is «very objective» and 1.0 is «very subjective». *Hatesonar* is a multi-class classifier, used to classify messages distinguishing between "hate speeches"/"offensive

²<https://textblob.readthedocs.io/en/dev/index.html>, retrieved on January 31st, 2022.

speeches”/“not offensive replies”. Each possibility is weighted with a confidence score (i.e., a probability score) within the range (0.0, 1.0) (Davidson, Warmsley, Macy, and Weber 2017).

In particular, 51 comments have been labelled “offensive” due to the presence of vulgar expressions, while 0 comments have been labelled “hate speech”. We also considered a second corpus of 1703 comments written in Italian, 44 of which have been labelled “offensive”, while only 1 has been labelled “hate speech”.

2. Hate speeches, fundamentalist discourses and conspiracies: problems of identification

The distinction between hate, vulgar, fundamentalist, and religious speech is difficult due to the known vagueness in the definition of “hate speech”. For example, such definitions as «any communication that disparages a person or a group on the basis of some characteristics (to be referred to as types of hate or hate classes) such as race, colour, ethnicity, gender, sexual orientation, nationality, religion, or other characteristics» (Nockleby 2000), focus mainly on the possible consequences of the meaning of the speech (which is, in semiotic terms, *substance of the content plane*, cf. Hjelmslev 1954) rather than on those elements and syntagms of the expression plane which manifest them. This leaves unanswered questions from the point of view of the Philosophy of Law: if the final decision on the criminal importance of a discourse is left to a judge, legal certainty is not ensured. If the definition of what is “hate” is vague, the speaker can’t know a priori whether his speech will be considered a crime or not.

Similar problems are related to conspiracy theories. Let us consider this repartee:

Pope Francis: The coronavirus has produced deaths and suffering, affecting the lives of all, especially that of the most vulnerable. Please don't forget the most vulnerable. #VaxLive #UnVaccinoperiPoveri

Answer: @Pontifex_it Governments subservient to world finance with the complicity of the “holy” father that cancelled experimental vaccines are making this life a hell on Earth

(Italian tweet translated by the author)

Does the message refer to a conspiracy? According to *TextBlob*, the reply scores 0.1 (polarity) and 0.4 (subjectivity). Thus, according to the classifier, it is a neutral, objective statement. According to *Hatesonar*, the answer is definitely not a hate speech (confidence: 0.04). The engine considers it slightly offensive, but its confidence is low (0.56). The probability that the message is neither a hate speech nor offensive is evaluated 0.3. Depending on our interpretation of the message, the “complicity” of the Holy Father can be only a criticism expressed toward the original message, not depicting a secret agreement between finance, politics, and religion. The expression “subservient to world finance” might not be that literal, being a cliché of radical discourse.

Recently, some MIT scholars proposed to employ machine learning techniques to train algorithms capable of detecting and identifying “language patterns” which features *fake news* (O’Brien, Latessa, Evangelopoulos, Boix 2018)³. In this perspective, *truth* ceases to be a correspondence between statements and a state of things, à la Wittgenstein (1922), and becomes the result of lexical and rhetoric choices. If the reply of the example above is submitted to the algorithm, the “fake news detector” labels the pattern «the world

³<http://fakenews.mit.edu/>

finance» as *most real pattern*, while «the “holy” father that cancelled», «experimental vaccines are making», and «hell on Earth» are labelled *most fake pattern*. In a similar way, starting from the (true) sentence as «a number of experimental vaccines are making their way through clinical testing», the engine labels *most fake pattern* the syntagms «experimental vaccines are making» and «clinical testing». Such kind of philosophically naïve approaches to the problem of truth a *meaning effect* (verisimilitude, plausibility) and truth, a relation that transcends language and implies metaphysical beliefs about “reality”. From this point of view, fake news, hate speeches, and conspiracy theories should be considered *literary genres*: paradoxically, it is possible to write “falsey” true news, i.e., true statements written in the typical style of fake news such as «3 Reasons Why You Should Stop Smoking Crack» or «Hilary admits it all: the so-called H₂O is only fresh water».

3. Hate speeches and semiosphere

The relation between hate speeches and culture is another important issue. The machine learning technique employed to train the neural networks of the aforementioned algorithms start from a corpus labelled by English speakers (usually, U.S. university students). This generates a bias, not only because there is a lack of classifiers trained in different languages and automatic translation of very large dataset is very imperfect, but also because the boundary between “illegal” and “admissible” is related to cultural specificities:

Italian (and Latin languages in general) are full of blasphemies and obscene expressions, while German is much more contained. Thus, an exclamation that in Italian would perhaps appear inappropriate, but not unusual, and would connote the speaker's origin and social level, in German would sound intolerably blasphemous or, in any case, exaggeratedly vulgar (Eco 2003).⁴

The social issues perceived in the considered culture can also be a bias. Tab. 1 displays two examples and their respective scores according to *TextBlob* and *Hatesonar*, while tab. 2 presents three “laboratory sentences” prepared by the author to test the algorithms.

<i>Text</i>	<i>Pol.</i>	<i>Subj.</i>	<i>Hate</i>	<i>Offensive</i>
A: @Pontifex Mary is dead. Do you think idolatry is the way to go in these end times? How bad do you want hell?	-0,45	0,5	0,08	0,65
B: @Pontifex_it But you like the Democratic Party and the shit of your friend Soros, LGBT, and various apocalyptic bullshit.	-0,1	0,65	0,09	0,89

Tab. 1: examples of answers and their score with *TextBlob* and *Hatesonar*

<i>Text</i>	<i>Pol.</i>	<i>Subj.</i>	<i>Hate</i>	<i>Offensive</i>
C: We are all part of a race. The human race	0,0	0,1	0,55	0,01
D: at least I am not a Jew	-0,3	0,4	0,46	0,13
E: at least I am not a Muslim	-0,3	0,4	0,04	0,36

Tab. 2: test sentences and their score with *TextBlob* and *Hatesonar*

⁴Eco’s consideration on blasphemy and translation are missing in the English version of the book. For this reason, the author of the present paper translated the passage from the original Italian text.

While there are no doubts on the reasons why the lexical choices which distinguish text B can be considered impolite, it is more difficult to understand why *Hatesonar* classifies the text A as offensive. It simply expresses religious fundamentalism and could have been written, for example, by an overzealous Protestant. As a simple commutation test reveals (Hjelmslev 1943), the reason why the message is considered vulgar by the algorithm lays in the term “hell”, which is a taboo subject in English, sometimes substituted with “heck”. *Hatesonar* is heavily biased by lexicon and underestimates the religious context, in which the term hell is not misplaced.

Reply B is interesting from the point of view of truth. Many political approaches to fake news or conspiracy theories challenge this kind of communication on the referential level. In this case, one would ask whether Soros is really responsible for the global financial crisis of the 1997, if he is really a puppet master secretly controlling the global economy and politics, if there is really such thing as a “gay lobby”, or if the Pope is a true liberal. The problem is that these messages do not ask the reader to evaluate them in terms of their truth value. They play an emotive and a conative function, even when they pretend to quote pseudo-scientific data on the presumed IQ of women or black people.

Similar problems affect the score of test sentence C. By the way, when the sentence is submitted to the MIT “fake news detector”, both «We are all» and «of a race. The human race» are labelled *most fake pattern*. Not amazingly, if sentence D is submitted, “at least I am not” is labelled *most real pattern*, while “Jew” is unlabelled. Once again, truth is not relevant to hate speeches and to their intended discriminatory effects.

The score of examples D and E is another case of lexical bias, which can be proved through a commutation test: while the term “Jew” raises the “hate speech” score, the use of the term “Muslim” has only a slightly higher probability to occur in offensive sentences, and example E is considered acceptable. Once again, the bias can be explained through the training process, and mirrors the issues and the contradictions of contemporary American society. In fact, in a corpus of hate speeches collected by researchers of the Berkeley D-Lab, the top 5 used words are Jews, White, Hate, Black, Women (cf. Galofaro, Toffano, Doan 2019). Thus, American culture perceives Muslims as a threat and Jews as the main target of racial hate. In Europe, in the Far East or in Russia the list and the rank of the issues would probably be very different.

4. Hashtags and replies

According to Guido Ferraro, the structure topic/comment, which he renames topic/*focus*, is universal and ancestral. It is the germ starting from which every complex narrative structure is unfolded. «The topic/focus model, of absolutely crucial importance, establishes the very possibility of languages» (Ferraro 2019: 288).

Leaving aside the limitations of the algorithms in detecting underlined above, a closer look to the corpus reveals that offensive replies focus only on specific topics and hashtags (tab. 3). Tweets present a topic/comment structure in which the topic can be expressed at the level of the expression plane by one or more hashtags, defining a *field of relevance*. The hashtag acts as a trigger stimulating answers.

<i>Topic manifested by the hashtag</i>	<i>Replies</i>	<i>Ratio of offensive replies (%)</i>
Israel and Palestine war	643	1,7
#LaudatoSiWeek (May 2021)	614	1,3
#LaudatoSi (period: 2015-2021)	215	0,9
#Myanmar	321	3,5
#Peace, #Fraternity (associated)	119	0,8

#Ascension	90	2,2
#WorldCommunicationDay	57	5,3
Child abuse	35	2,9
#PrayTogether, #DayofFamilies (associated)	18	5,6
#Birthrate	11	18

Tab. 3: offensive replies per topic/hashtag

As we see, the ratio of offensive replies significantly increases when the answers are a small number. Apparently, a loyal patrol of “Pope-trolls” attends Pope Francis’ twitter account for the sole purpose to argue. It seems an interesting case of reverse-retention of anti-followers, a sort of “anti-evangelization”. Except from that, ethical issues attract the attention of militant atheists but also of conservative Catholics. The number of the latter group has grown over the years as the Pope’ pro-communist reputation, as it is demonstrated by the difference between the hashtags #LaudatoSi and #LaudatoSiWeek. In our corpus, the average subjectivity score of the offensive and non-offensive tweets is the same.

Finally, approximatively 500 tweets are formulas such as “amen”, “hear us” etc. This confirms the hypothesis of the “trigger”: the structure tweet/reply is similar to the prayer of the faithful in Catholic liturgy of the Word (e.g., «Inspire all the baptized and confirmed to be faithful to their promises. Fill them with zeal and courage» - «Lord, hear your people’s prayer»):

Pope Francis: Consistent #prayer produces progressive transformation, makes us strong in times of tribulation, gives us the grace to be supported by Him who loves us and always protects us.

Retweet: Prayer is the key of our life amen

Interestingly, some offensive replies follow the same syntagmatic structure, substituting “amen” with such insults as “fuck your prayers”, “go to hell”, “eat shit” etc.

5. A conflict of isotopies

This section will investigate the semantic consequences of the reply in relation to the meaning of the original message. The analysis technique has been presented in Galofaro, Toffano, Doan (2019). To represent the semantic space of the document, the algorithm makes use of quantum information retrieval on the contexts created by the co-occurrence of the lexical elements of the message. Each “dissemination” of a lexeme through the different contexts of the document will be considered a word-vector intersecting all the other words, manifesting an *isotopy* (Greimas and Courtés 1979):

Isotopy at first designated iterativity along a syntagmatic chain of classemes which assure the homogeneity of the utterance-discourse. From this point of view, it is clear that the syntagm joining together at least two semic figures may be considered the minimal context necessary for establishing an isotopy (Greimas and Courtés 1979: 163).

In particular, the algorithm measures correlation and entanglement between word vectors, which we can interpret in terms of isotopic proximity/distance in the semantic space of the document, and allows a comparison between the scores of the original tweet and of the new text composed by both the tweet and the answer. In general, while

the correlation/anticorrelation score provides information on the proximity of two isotopies, the entanglement score measures the non-separability of two lexemes into different isotopies. In other terms, in the considered text, entangled lexemes give life to one and only one isotopy from a semantic point of view. This technique grants the machine access to information about the form of the content plane of the message, i.e. about its meaning. The model will be tested on the following text:

Pope Francis: On this Feast of the #Ascension, while we contemplate Heaven, where Christ has ascended and sits at the right hand of the Father, let us ask Mary, Queen of Heaven, to help us to be courageous witnesses to the Risen One in the world, in the concrete situations of life.

Reply: @Pontifex: Mary's dead. Do you think idolatry is the way to go in these end times? How bad do you want hell?

The semantic space of the original tweet and of the new document composed by the tweet and the reply is represented in fig. 1.

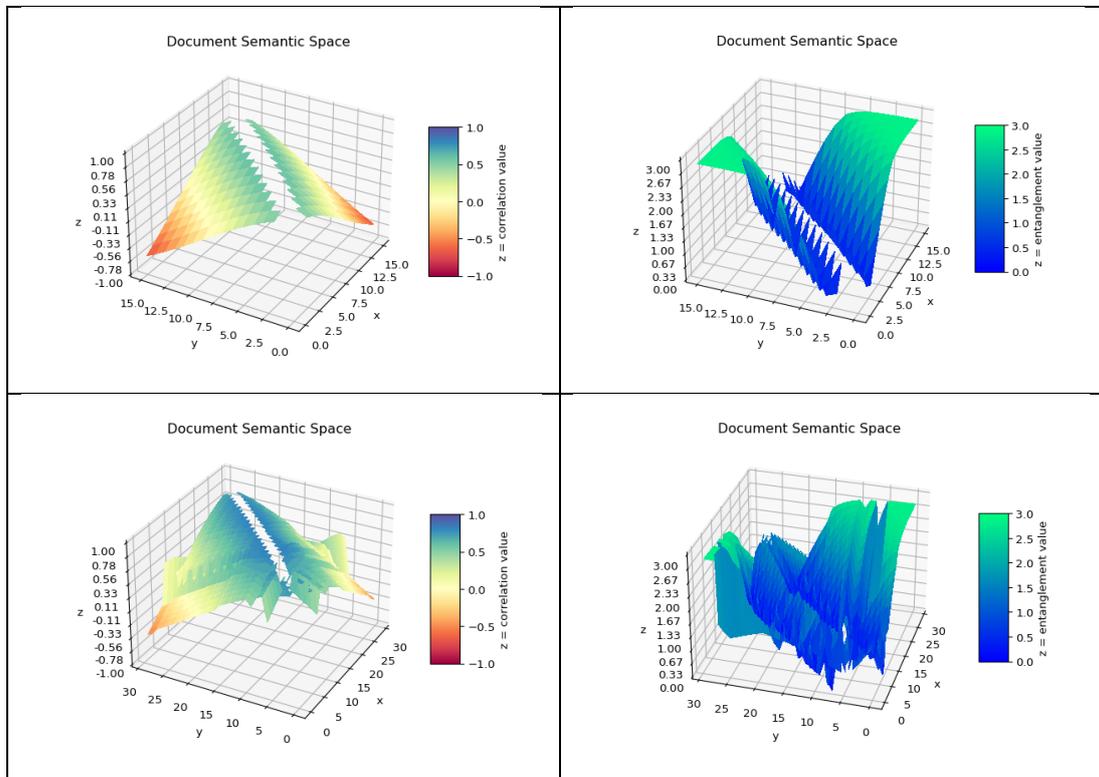


Fig. 1: above: the semantic space of the original tweet in terms of correlation (left) and entanglement (right) between lexemes. Below: the semantic space of the original tweet is modified by the reply.

What follows is a technical presentation of the model. The not interested reader can directly read the next paragraph. In fig. 1, the numbers on the x and y axis are indexes of the word-vectors (lexemes) composing the document. Each word-vector encodes the distance with all the other word-vectors of the document. The z-axis represents the degree of anticorrelation-correlation, within the range (-1.0, + 1.0), or entanglement, within the range (0, + 2,8). Regarding correlation, the maximal values of -1 represents the maximum of anticorrelation between lexemes (*allotopy*); the value of +1 represents

the maximum of correlation; the value of 0 represent the absence of correlation between the two. Regarding entanglement, values above 2 represents a non-classic (i.e. quantum) correlation: thus, the two lexemes belongs to the same semantic isotopy. A detailed presentation of the model can be found in Galofaro, Toffano, Doan (2019).

In the text above, a particular isotopy is represented by the hashtag #Ascension. According to the hypothesis presented in the previous section, the hashtag manifests the *topic* of the tweet. Fig. 2 represents the section of the semantic space of the document corresponding to the lexeme #Ascension (and, of course, “ascended”).

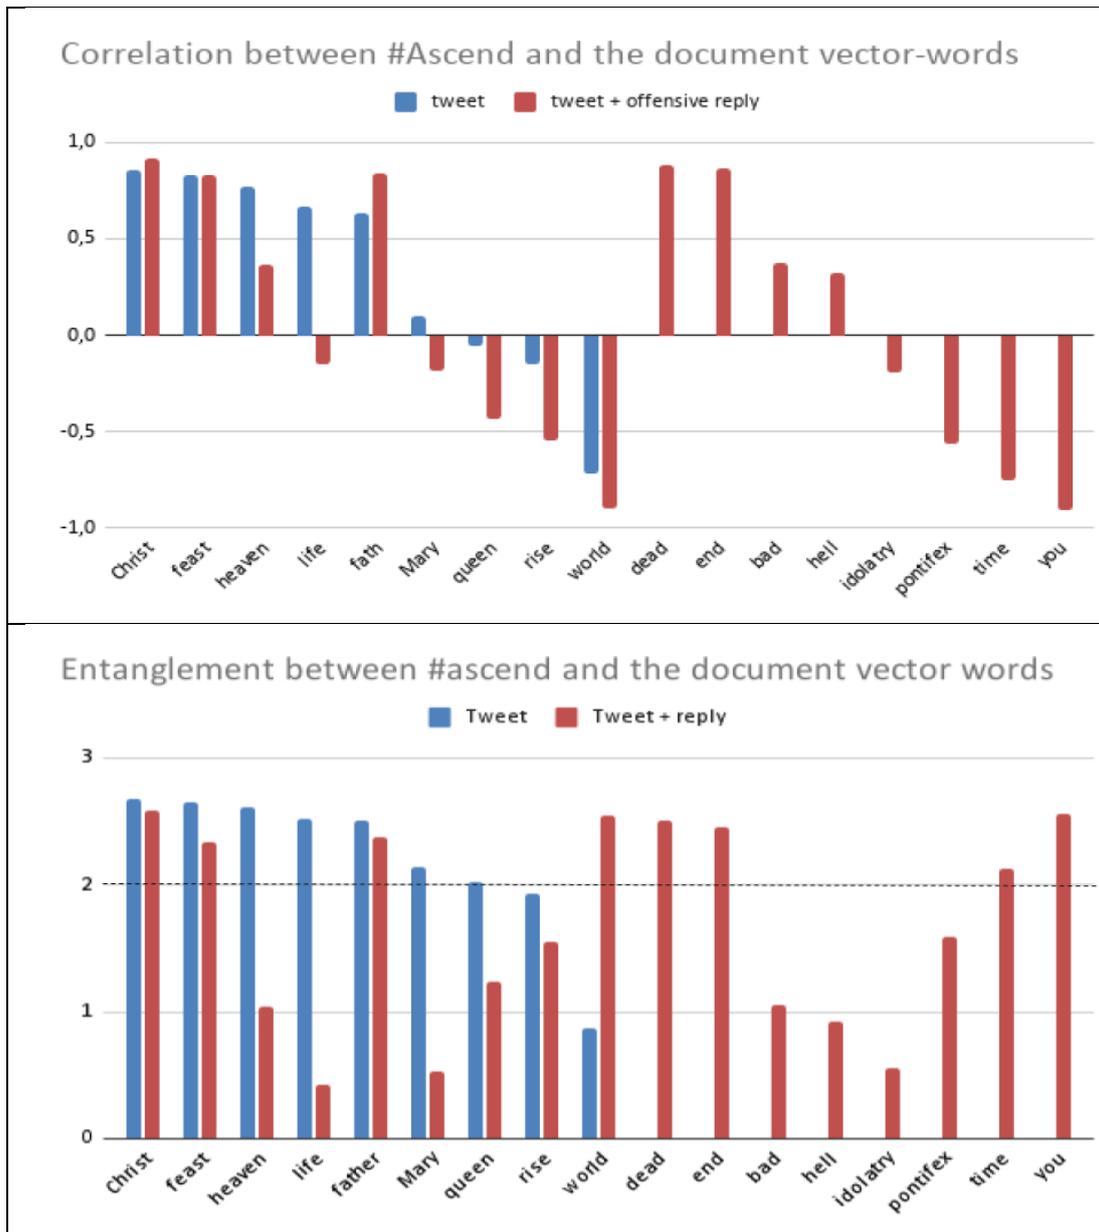
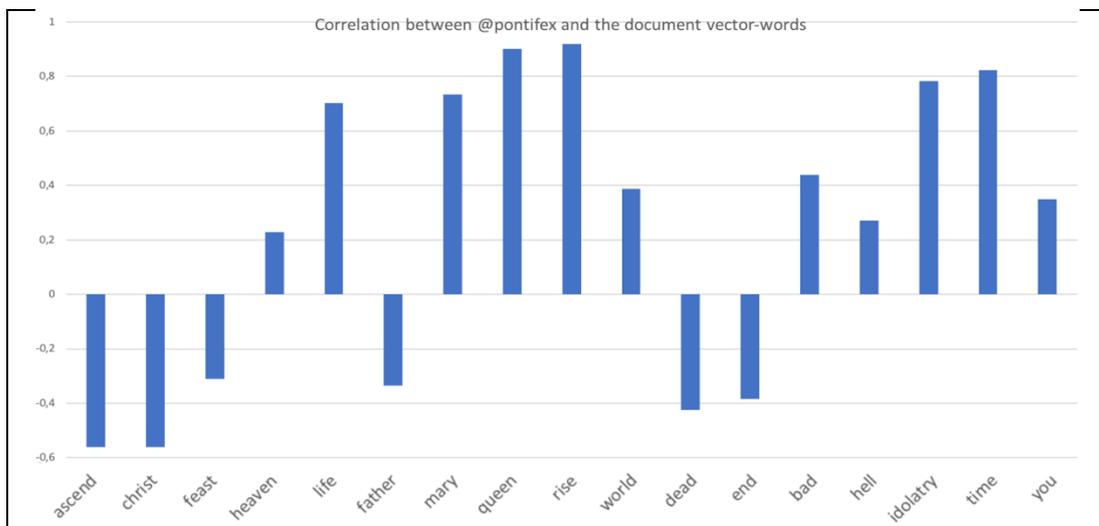


Fig. 2: Values of the word-vector #ascend in correspondence to the other word-vectors. It is possible to see how the reply modifies the original isotopy relating it to new lexemes and creating an opposition with other ones. Interestingly, some terms entangled in the original message are not entangled in the second one (heaven, life, Mary), while new entanglement relations are created.

A closer look to fig. 1 let emerge how, in Pope Francis' original tweet, the topic #ascend is isotopically connected to /Christ/, /feast/, /Heaven/, /life/, /Father/, and, weakly, to /Mary/. Furthermore, /ascend/ is strongly anticorrelated to /world/. But, when modified by the second tweet, these relations are reshaped. In particular, #Ascension is no more entangled with /Heaven/, /life/, and /Mary/: they are not part of the same isotopy anymore and their meaning changes. The reply opposes death to life, the value which features the original tweet. #Ascension is now entangled with "world", "dead", "end", and "you". Furthermore, now #Ascension is anticorrelated to life. The model seems adequate to the perceived meaning of the document. In fact, the Pope proposes a homologation: Heaven/World = Christ/Mary = Top/Bottom = Contemplation/Witness. The reply selects /Mary/ to reshape the homologation, associating it a new term: life/death, and linking the latter to /Mary/, /you/, /end/, /times/. Since /Mary/ was connected to #Ascension in the original document, all the previous isotopies are modified and connected to the new ones: death features now human world.

6. The address sign @

As well as the hashtag, the address sign @ can be considered a morphologic mark. The receiver of the tweet is represented via it in the morphology of the language. It raises the spectre of the reader, for it is not possible to know whether the receiver will really read the message unless he replies and, even in this case, the writer of the answer could not be the "real" addresser (e.g. the Pope), but a ghost writer, a press officer, a replacement of sort. Provided that, the address sign constructs the meaning effect of the reader's presence, summoning him. Since the social network implies a "public" reader, the sign @ should be considered a special case of *apostrophe*, a rhetorical device which occurs when the writer «speaks directly and vividly to a person other than the natural or conventional recipient of the speech itself» (Mortara Garavelli 1988: 268). Regarding enunciation, the apostrophe implies the insertion of the marks of the *personal enunciation* in an overall impersonal narration. The effect is relative to the *conative* and *emotive function* of communication. Both *execration* and *curse* are special genres of apostrophe. «The referent can indifferently be animate or inanimate, alive or dead, real or imaginary» (Beccaria 1994: 75). Thus, the fact that the addressee replies or not it is irrelevant to the effectiveness of the rhetoric device.



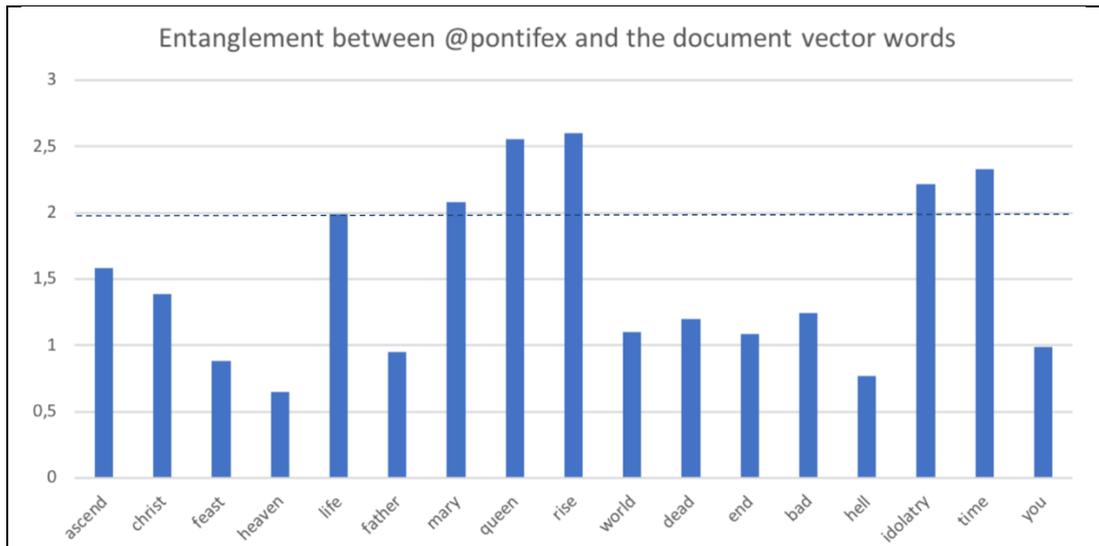


Fig. 2: Values of the word-vector @Pontifex in correspondence to the other word-vectors.

Looking at fig. 2, it is possible to see the value scored by the isotopic vector /@Pontifex/ when intersecting the other vectors in the semantic space generated by the original tweet and the reply. In this case one cannot appreciate the “reshaping” of the meaning, since /pontifex/ is present only in the answer. However, we can see how /@Pontifex/ is correlated and entangled to form a non-separable system with the vectors /Mary/, /Queen/, /rise/, /idolatry/ and /time/ and strongly anticorrelated to /ascension/ and /Christ/. This seems adequate to the perceived meaning of the document: At the end of time, the Pontifex represents the idolatrous cult of Mary, which is opposed to the “true” worship of Christ. Since the sign @ directly addresses these negative values to the Pontifex, the meaning effect constructed by the message belongs to the rhetoric subgenre of the execration. Interestingly, /@Pontifex/ is not entangled and only weakly related to /you/. This is a limit of the algorithm, which does not make inferences about the reference shared by the two terms since it evaluates only the contexts and has no notions on the indexical function played by the pronoun /you/.

7. Discussion

Unlike the classifiers trained through machine learning techniques, quantum structural semantics clarifies the relationship between tweets and offensive answers and the way in which the second, triggered by the topic of the first, challenges its meaning by reshaping it making use of the rhetoric device of the apostrophe to modify the enunciation of the message. While the considered machine learning-based classifiers are biased by an excessive weight of the lexical component of the messages, quantum structural semantics allows to consider the context in which lexemes occur. Furthermore, quantum structural semantics does not imply a training process involving humans. Thus, it is not biased by the social and cultural issues which features the semiosphere in which classifiers are trained. Hopefully, in the future quantum structural semantics will provide new insights on natural language processing, avoiding naïve and empiricist philosophical views on truth, objectivity, and law which currently flaw many machine learning-based researches

References

- Beccaria, Gian Luigi (1994), *Dizionario di linguistica e di filologia, metrica, retorica*, Einaudi, Torino 2004.
- Ceriani, Giulia (2015), *Reputazione, popolarità, sentiment: discorsi sul sesso in rete e problematiche di ricerca integrate*, in Ferraro Guido e Lorusso Anna Maria, a cura di, *Nuove forme di interazione dal Web al Mobile*, Libellula edizioni, Lecce, pp. 53-52.
- Davidson, Thomas, Warmesley, Dana, Macy, Michael W., and Weber, Ingmar (2017), «Automated Hate Speech Detection and the Problem of Offensive Language», in *ICWSM*, <https://arxiv.org/pdf/1703.04009.pdf>
- Eco, Umberto (2003), *Dire quasi la stessa cosa: esperienze di traduzione*, Bompiani, Milano 2013 (ebook version).
- Ferraro, Guido (2019), *Semiotica 3.0: 50 idee-chiave per un rilancio della scienza della significazione*, Aracne, Roma.
- Galofaro, Francesco, Toffano, Zeno, Doan, and Bich-Liên (2019), «Semantic Quantum Correlations in Hate Speeches», in *Rivista Italiana di Filosofia del Linguaggio*, SFL - Language and emotions, pp. 370-383.
- Greimas, Algirdas J., and Courtés, Joseph (1979), *Sémiotique: dictionnaire raisonné de la théorie du langage*, Hachette, Paris (*Semiotics and Language: An Analytical Dictionary*, Indiana University Press, Bloomington 1982).
- Hjelmslev, Louis (1943), *Omkring sprogteoriens grundlaeggelse*, Akademisk forlag, Copenhagen (*Prolegomena to a theory of language*, transl. by Francis Madison Whitfield, University of Wisconsin, 1969).
- Hjelmslev, Louis (1954), «La stratification du langage», in *Word*, 10, pp. 163-188.
- Melucci, Massimo (2015), *Introduction to Information Retrieval and Quantum Mechanics*, Springer, Berlin.
- Mortara Garavelli, Bice (1988), *Manuale di retorica*, Bompiani, Milano 2003.
- Nockleby, John T. (2000), *Hate Speech*, in Leonard W. Levy, Kenneth L. Karst et al., *Encyclopedia of the American Constitution*, Macmillan, New York, pp. 1277-1279.
- Peverini, Paolo, Lorusso, Anna Maria (2017), *Il racconto di Francesco: la comunicazione del Papa nell'era della connessione globale*, Luiss University Press, Roma.
- Santangelo, Antonio (2015), *Twitter e l'analisi semiotica di un programma televisivo*, in Guido Ferraro e Anna Maria Lorusso (2015), *Nuove forme di interazione dal Web al Mobile*, Libellula edizioni, Lecce, pp. 67-88.
- Van Rijsbergen, Cornelis J. (2004), *The Geometry of Information Retrieval*, Cambridge University Press, Cambridge UK.
- Wittgenstein, Ludwig (1922), *Tractatus Logico-Philosophicus*, Kegan Paul, ed. by, Basil Blackwell, Oxford 1961.