

ItEM: A Vector Space Model to Bootstrap an Italian Emotive Lexicon

Lucia C. Passaro, Laura Pollacci, Alessandro Lenci
ColingLab, Dipartimento di Filologia, Letteratura e Linguistica
University of Pisa (Italy)

lucia.passaro@for.unipi.it, laurapollacci.pl@gmail.com,
alessandro.lenci@unipi.it

Abstract

English. In recent years computational linguistics has seen a rising interest in subjectivity, opinions, feelings and emotions. Even though great attention has been given to polarity recognition, the research in emotion detection has had to rely on small emotion resources. In this paper, we present a methodology to build emotive lexicons by jointly exploiting vector space models and human annotation, and we provide the first results of the evaluation with a crowdsourcing experiment.

Italiano. Negli ultimi anni si è affermato un crescente interesse per soggettività, opinioni e sentimenti. Nonostante sia stato dato molto spazio al riconoscimento della polarità, esistono ancora poche risorse disponibili per il riconoscimento di emozioni. In questo lavoro presentiamo una metodologia per la creazione di un lessico emotivo, sfruttando annotazione manuale e spazi distribuzionali, e forniamo i primi risultati della valutazione effettuata tramite crowdsourcing.

1 Introduction and related work

In recent years, computational linguistics has seen a rising interest in subjectivity, opinions, feelings and emotions. Such a new trend is leading to the development of novel methods to automatically classify the emotions expressed in an opinionated piece of text (for an overview, see Liu, 2012; Pang and Lee, 2008), as well as to the building of annotated lexical resources like SentiWordNet (Esuli and Sebastiani, 2006; Das and Bandyopadhyay, 2010), WordNet Affect (Strapparava and Valitutti, 2004) or EmoLex (Mohammad and Turney, 2013). Emotion detection can be useful in several applications, e.g. in Customer Relationship Management (CRM) it can be used to track sentiments towards companies and their services, products or others target enti-

ties. Another kind of application is in Government Intelligence, to collect people's emotions and points of views about government decisions. The common trait of most of these approaches is a binary categorization of emotions, articulated along the key opposition between POSITIVE and NEGATIVE emotions. Typically, then, these systems would associate words like "rain" and "betray" to the same emotion class in that they both evoke negative emotions, without further distinguishing between the SADNESS-evoking nature of the former and the ANGER-evoking nature of the latter. Emotion lexica, in which lemmas are associated to the emotions they evoke, are valuable resources that can help the development of detection algorithms, for instance as knowledge sources for the building of statistical models and as gold standards for the comparison of existing approaches. Almost all languages but English lack a high-coverage high-quality emotion inventory of this sort. Building these resources is very costly and requires a lot of manual effort by human annotators. On the other hand, connotation is a cultural phenomenon that may vary greatly between languages and between different time spans (Das and Bandyopadhyay, 2010), so that the simple transfer of an emotive lexicon from another language cannot be seen as nothing else than a temporary solution for research purposes.

Crowdsourcing is usually able to speed the process and dramatically lower the cost of human annotation (Snow et al., 2008; Munro et al, 2010). Mohammad and Turney (2010, 2013) show how the "wisdom of the crowds" can be effectively exploited to build a lexicon of emotion associations for more than 24,200 word senses. For the creation of their lexicon, EmoLex, they selected the terms from Macquarie Thesaurus (Bernard, 1986), General Inquirer (Stone et al., 1966), WordNet Affect Lexicon (Strapparava and Valitutti., 2004) and Google n-gram corpus (Brants and Franz, 2006) and they exploited a crowdsourcing experiment, in order to obtain, for every target term, an indication of

its polarity and of its association with one of the eight Plutchik (1994)’s basic emotions (see below). The methodology proposed by Mohammad and Turney (2010, 2013), however, cannot be easily exported to languages where even small emotive lexica are missing. Moreover, a potential problem of a lexicon built solely on crowdsourcing techniques is that its update requires a re-annotation process. In this work we’re proposing an approach to address these issues by jointly exploiting corpus-based methods and human annotation. Our output is ItEM, a high-coverage emotion lexicon for Italian, in which each target term is provided of an association score with eight basic emotion. Given the way it is built, ItEM is not only a static lexicon, since it also provides a dynamic method to continuously update the emotion value of words, as well as to increment its coverage. This resource will be comparable in size to EmoLex, with the following advantages: i) minimal use of external resources to collect the seed terms; ii) little annotation work is required to build the lexicon; iii) its update is mostly automatized.

This paper is structured as follows: In section 2, we present ItEM by describing its approach to the seed collection and annotation step, its distributional expansion and its validation. Section 3 reports the results obtained from the validation of the resource using a crowdsourcing experiment.

2 ItEM

Following the approach in Mohammad and Turney (2010, 2013), we borrow our emotions inventory from Plutchik (1994), who distinguishes eight “basic” human emotions: JOY, SADNESS, ANGER, FEAR, TRUST, DISGUST, SURPRISE and ANTICIPATION. Positive characteristics of this classification include the relative low number of distinctions encoded, as well as its being balanced with respect to positive and negative feelings. For instance, an emotive lexicon implementing the Plutchik’s taxonomy will encode words like “ridere” (laugh) or “festa” (celebration) as highly associated to JOY while words like “rain” (pioggia) or “povertà” (poverty) will be associated to SADNESS, and words like “rissa” (fight) or “tradimento” (betray) will be encoded as ANGER-evoking entries.

ItEM has been built with a three stage process: In the first phase, we used an online feature elicitation paradigm to collect and annotate a small set of emotional seed lemmas. In a second phase,

we exploited distributional semantic methods to expand these seeds and populate ItEM. Finally, our automatically extracted emotive annotations have been evaluated with crowdsourcing.

2.1 Seed collection and annotation

The goal of the first phase is to collect a small lexicon of “emotive lemmas”, highly associated to the one or more Plutchik’s basic emotions. To address this issue, we used an online feature elicitation paradigm, in which 60 Italian native speakers of different age groups, levels of education, and backgrounds were asked to list, for each of our eight basic emotions, 5 lemmas for each of our Parts-of-Speech (PoS) of interest (Nouns, Adjectives and Verbs). In this way, we collected a lexicon of 347 lemmas strongly associated with one or more Plutchik’s emotions. For each lemma, we calculated its emotion distinctiveness as the production frequency of the lemma (i.e. the numbers of subjects that produced it) divided by the number of the emotions for which the lemma was generated. In order to select the best set of seed to use in the bootstrapping step, we only selected from ItEM the terms evoked by a single emotion, having a distinctiveness score equal to 1. In addition, we expanded this set of seeds with the names of the emotions such as the nouns “gioia” (joy) or “rabbia” (anger) and their synonyms attested in WordNet (Fellbaum, 1998), WordNet Affect (Strapparava and Valitutti, 2004) and Treccani Online Dictionary (www.treccani.it/vocabolario).

Emotion	N. of seeds	Adj	Nouns	Verbs
Joy	61	19	26	19
Anger	77	32	30	16
Surprise	60	25	17	22
Disgust	80	40	21	25
Fear	78	37	20	27
Sadness	77	39	22	26
Trust	62	25	21	17
Anticipation	60	15	22	23

Table 1 Distribution of the seeds lemmas

Globally, we selected 555 emotive seeds, whose distribution towards emotion and PoS is described in Table 1.

2.2 Bootstrapping ItEM

The seed lemmas collected in the first phase have been used to bootstrap ItEM using a corpus-based model inspired to Turney and Littmann (2003) to automatically infer the semantic orientation of a word from its distributional similarity with a set of positive and negative paradigm

words. Even if we employ a bigger number of emotion classes, our model is based on the same assumption that, in a vector space model (Sahlgren, 2006; Pantel and Turney, 2010), words tend to share the same connotation of their neighbours. We extracted from the La Repubblica corpus (Baroni et al, 2004) and itWaC (Baroni et al., 2009), the list T of the 30,000 most frequent nouns, verbs and adjectives, which were used as target and contexts in a matrix of co-occurrences extracted within a five word window (± 2 words, centered on the target lemma, before removing the words not belonging to T). Differently from the Turney and Littmann (2003)’s proposal, however, we did not calculate our scores by computing the similarity of each new vector against the whole sets of seed terms. On the contrary, for each $\langle emotion, PoS \rangle$ pair we built a centroid vector from the vectors of the seeds belonging to that emotion and PoS, obtaining in total 24 centroids. We constructed different word spaces according to PoS because the context that best captures the meaning of a word, differs depending on the word to be represented (Rothenhäusler and Schütze, 2007). Finally, our emotionality scores have been calculated on the basis of the distance between the new lemmas and the centroid vectors. In this way, each target term received a score for each basic emotion. In order to build the vector space model, we re-weighted the co-occurrence matrix using the Positive Pointwise Mutual Information (Church and Hanks, 1990), which works well for both word–context matrices (Pantel & Lin, 2002a) and term–document matrices (Pantel & Lin, 2002b). In particular, we used the Positive PMI (PPMI), in which negative scores are changed to zero, and only positive ones are considered (Niwa & Nitta, 1994). We followed the approach in Polajnar and Clark (2014) by selecting the top 240 contexts for each target word. Finally, we calculated the emotive score for a target word as the cosine similarity with the corresponding centroid (e.g. the centroid of “JOY-nouns”).

The output of this stage is a list of words ranked according to their emotive score. Appendix A shows the top most associated adjectives, nouns and verbs in ItEM. As expected, a lot of target words have a high association score with more than one emotive class, and therefore some centroids are less discriminating because they have a similar distributional profile. Figure 1 shows the cosine similarity between the emotive centroids: we can observe, for example, a high similarity between SADNESS and FEAR, as well as between

SURPRISE and JOY. This is consistent with the close relatedness between these emotions.

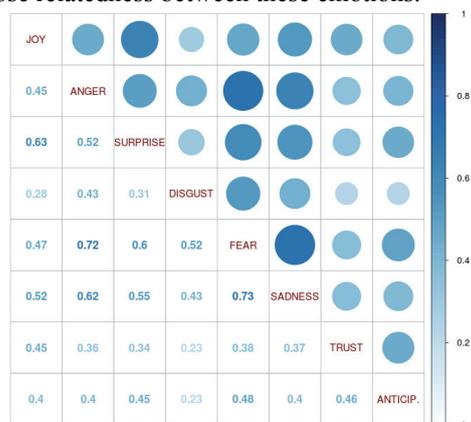


Figure 1 Cosine similarity between the emotive centroids

2.3 Validation

We evaluated our procedure using a two-step crowdsourcing approach: in the first step, for each $\langle emotion, PoS \rangle$ pair we ranked the target words with respect to their cosine similarity with the corresponding emotive centroid. We then selected the top 50 words for each centroid and we asked the annotators to provide an emotive score: Given a target word $\langle w \rangle$, for each Plutchik’s emotion $\langle e \rangle$, three annotators were asked to answer the question “How much is $\langle w \rangle$ associated with the emotion $\langle e \rangle$?”. The annotators had to choose a score ranging from 1 (not associated) to 5 (highly associated). Since very often the words may be associated with more than one emotion, we wanted to estimate the average degree of association between the word and the various emotions. Empirically, we defined the best distinctiveness score d as follows:

$$d = \frac{max1 - max2}{me * (max1 - mn)}$$

Where $max1$ is the highest emotive association for the target word, $max2$ is the second higher value, me is the association score between the target word and the target emotion, and mn is the average of the evaluations for the word across the emotion classes. This formula captures the intuition that a word is distinctive for a target emotion if its association degree with the target emotion is high as well as its association degree with the other classes is low. After ranking the words over this association score, we selected the top 10 distinctive nouns, adjectives and verbs for each $\langle emotion, PoS \rangle$ pair, in order to further expand the set of the seeds used to build the distributional space. For this second run, we removed the words belonging the top 10 of more than one emotion and we added the remaining

192 words to the set of the seeds used to build the centroid emotive vectors, using the procedure described in section 2.2. The second run allows us to evaluate the quality of the initial seeds and to discover new highly emotive words.

3 Results

We have evaluated the precision of our distributional method to find words correctly associated with a given emotion, as well the effect of the incremental process of seed expansion. In particular, we evaluated the top 50 nouns, adjectives and verbs for each emotion. Precision@50 has been calculated by comparing the vector space model’s candidates against the annotation obtained with crowdsourcing. True positives (TP) are the words found in the top 50 neighbours for a particular emotion and PoS, for which the annotators provided a average association score greater than 3. False positives (FP) are the words found in the top 50 nouns, adjectives and verbs, but for which the aggregate evaluation of the evaluators is equal or lower than 3. Table 2 shows the Precision by emotion in the first run (P Run 1) and in the second one (P Run 2), calculated on a total of 1,200 target associations.

Emotion	P (Run 1)	P (Run 2)
Joy	0.787	0.767
Anger	0.813	0.827
Surprise	0.573	0.56
Disgust	0.78	0.753
Fear	0.673	0.727
Sadness	0.827	0.793
Trust	0.43	0.5
Anticipation	0.557	0.527
Micro AVG	0.68	0.682

Table 2 Precision by Emotion (Runs 1 and 2)

If we analyze the same results by aggregating the Precision by PoS (Table 3), we can notice some differences between the first and the second run. Although overall there is a slight increase of the Precision score, this growth only affects verbs and adjectives. This is probably due to the way in which the noun seeds are distributed around the emotion centroids: a lot of them, in fact, are strongly associated to more than one emotion.

PoS	P (Run 1)	P (Run 2)
Adjectives	0.727	0.735
Nouns	0.685	0.675
Verbs	0.629	0.635

Table 3 Precision by PoS (Runs 1 and 2)

To appreciate the gain obtained in the second run, we analyzed the medium change of cosine similarity between the first and the second ex-

periment, and we noticed that the true positives have, on average, a higher cosine similarity with the corresponding emotive centroid in the second run (cf. Table 4). This proves the positive effect produced by the new seeds discovered by the distributional model in the first run.

Emotion	CosR1	CosR2	CosR2-CosR1
Joy	0.564	0.595	+0.032
Anger	0.582	0.6	+0.018
Surprise	0.635	0.657	+0.022
Disgust	0.524	0.555	+0.034
Fear	0.616	0.613	-0.003
Sadness	0.612	0.648	+0.036
Trust	0.575	0.665	+0.103
Anticipation	0.54	0.563	+0.027
Macro Avg	0.581	0.612	+0.034

Table 4 Increase of cosine similarity

In general, the distributional method is able to achieve very high levels of precision, despite an important variance among emotion types. Some of them (e.g., ANTICIPATION) confirm to be quite hard, possibly due to a higher degree of vagueness in their definition that might also affect the intuition of the evaluators.

The results that we achieved for the different emotions and PoS show that additional research is needed to improve the seed selection phase, as well as the tuning of the distributional space.

4 Conclusion

What we are proposing with ItEM is a reliable methodology that can be very useful for languages that lack lexical resources for emotion detection, and that is at the same time scalable and reliable. Moreover, the resulting resource can be easily updated by means of fully automatic corpus-based algorithms that do not require further work by human annotators, a vantage that can turn out to be crucial in the study of a very unstable phenomenon like emotional connotation.

The results of the evaluation with crowdsourcing show that a seed-based distributional semantic model is able to produce high quality emotion scores for the target words, which can also be used to dynamically expand and refine the emotion tagging process.

Reference

- Baroni M. and Lenci A. (2010). Distributional Memory: A General Framework for Corpus-Based Semantics. In *Computational Linguistics*, 36 (4), pp. 673-721.
- Baroni M., Bernardini S., Comastri F., Piccioni L., Volpi A., Aston G. and Mazzoleni M. (2004). Introducing the “la Repubblica” Corpus: A Large, Annotated, TEI(XML)-Compliant Corpus of Newspaper Italian. In *Proceedings of LREC 2004*.
- Baroni M., Bernardini S., Ferraresi A. and Zanchetta E. (2009). The WaCky Wide Web: A Collection of Very Large Linguistically Processed Web-Crawled Corpora. *Language Resources and Evaluation* 43(3): 209-22
- Bradley, M. and Lang P. (1999) Affective norms for english words (ANEW): Instruction manual and affective ratings. Technical Report, C-1, The Center for Research in Psychophysiology, University of Florida.
- Brants T. and Franz A. (2006). Web 1t 5-gram version 1. Linguistic Data Consortium.
- Das A. and Bandyopadhyay S. (2010). Towards the Global SentiWordNet. In *Proceedings of PACLIC 2010*, pp. 799-808.
- Esuli A. and Sebastiani F. (2006). SentiWordNet: A publicly available lexical resource for opinion mining. In *Proceedings of LREC 2006*.
- Fellbaum C. (1998). *WordNet: An Electronic Lexical Database*. Cambridge, MA: MIT Press
- Liu B. (2012). *Sentiment Analysis and Opinion Mining*. Morgan & Claypool Publishers.
- Mohammad S. M. and Turney P. D. (2010). Emotions Evoked by Common Words and Phrases: Using Mechanical Turk to Create an Emotion Lexicon. In *Proceedings of the NAACL-HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pp. 26-34.
- Mohammad S. M. and Turney P. D. (2013). Crowdsourcing a Word-Emotion Association Lexicon. In *Computational Intelligence*, 29 (3), pp. 436-465.
- Munro R., Bethard S., Kuperman V., Lai V., Melnick R., Potts C., Schnoebelenand T., Tily H. (2010). Crowdsourcing and language studies: the new generation of linguistic data. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*.
- Niwa, Y., & Nitta, Y. (1994). Co-occurrence vectors from corpora vs. distance vectors from dictionaries. In *Proceedings of the 15th International Conference On Computational Linguistics*, pp. 304-309, Kyoto, Japan.
- Pang B. and Lee L. (2008). Opinion mining and sentiment analysis. In *Foundations and trends in Information Retrieval*, 2 (1-2), pp. 1-135.
- Pantel, P., & Lin, D. (2002a). Discovering word senses from text. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 613–619, Edmonton, Canada.
- Pantel, P., & Lin, D. (2002b). Document clustering with committees. In *Proceedings of the 25th Annual International ACM SIGIR Conference*, pp. 199–206.
- Plutchik R. (1994) *The psychology and biology of emotion*. Harper Collins, New York.
- Polajnar T., and Clark S. (2014). Improving Distributional Semantic Vectors through Context Selection and Normalisation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2014)*, pp. 230-238, Gothenburg, Sweden, 2014.
- Rothenhäusler, K. and Schütze, H. (2007) Part of Speech Filtered Word Spaces. In *Proceedings of the Sixth International and Interdisciplinary Conference on Modeling and Using Context*, Roskilde, Denmark, August 20-24, 2007.
- Sahlgren M. (2006) *The Word-Space Model: Using distributional analysis to represent syntagmatic and paradigmatic relations between words in high-dimensional vector spaces*. Ph.D. dissertation, Stockholm University.
- Snow R., O'Connor, B. Jurafsky D, Ng, A. (2008) Cheap and fast - but is it good? Evaluating non-expert annotations for natural language tasks. In *Proceedings of EMNLP 2008*, pp. 254-263.
- Stone P., Dunphy D.C., Smith M.S. and Ogilvie D.M. (1966). *The General Inquirer: A Computer Approach to Content Analysis*. The MIT Press, Cambridge, MA.
- Strapparava C. and Valitutti A. (2004) Wordnet-Affect: An affective extension of WordNet. *Proceedings of LREC-2004*, pp. 1083-1086.
- Turney P. D and Pantel P. (2010). From frequency to meaning: vector space models for semantics. In *Journal of Artificial Intelligence Research*, 37, pp. 141-188.
- Turney P.D. and Littman M.L. (2003) Measuring praise and criticism: Inference of semantic orientation from association. In *ACM Transactions on Information Systems*.

Appendix A: Top 5 adjectives, verbs and nouns for each emotion, with their association scores, calculated as the cosine similarity between the word and the corresponding centroid vector.

EMOTION	ADJECTIVES	COSINE	VERBS	COSINE	NOUNS	COSINE
JOY	gioioso (<i>joyful</i>)	0.85	rallegrare (<i>to make happy</i>)	0.6	gioia (<i>joy</i>)	0.83
	scanzonato (<i>easygoing</i>)	0.68	consolare (<i>to comfort</i>)	0.54	ilarità (<i>cheerfulness</i>)	0.73
	spiritoso (<i>funny</i>)	0.66	apprezzare (<i>to appraise</i>)	0.53	tenerenza (<i>tenderness</i>)	0.72
	scherzoso (<i>joking</i>)	0.65	applaudire (<i>to applaud</i>)	0.53	meraviglia (<i>astonishment</i>)	0.7
	disinvolto (<i>relaxed</i>)	0.62	rammentare (<i>to remind</i>)	0.53	commozione (<i>deep feeling</i>)	0.69
ANGER	insofferente (<i>intolerant</i>)	0.72	inveire (<i>to inveigh</i>)	0.59	impazienza (<i>impatience</i>)	0.8
	impaziente (<i>anxious</i>)	0.67	maltrattare (<i>totreatbadly</i>)	0.58	dispetto (<i>prank</i>)	0.76
	permaloso (<i>prickly</i>)	0.66	offendere (<i>to offend</i>)	0.56	rancore (<i>resentment</i>)	0.75
	geloso (<i>jealous</i>)	0.66	ingiuriare (<i>to vituperate</i>)	0.53	insofferenza (<i>intolerance</i>)	0.74
	antipatico (<i>unpleasant</i>)	0.65	bastonare (<i>to beat with a cane</i>)	0.52	antipatia (<i>impatience</i>)	0.74
SURPRISE	perplesso (<i>perplexed</i>)	0.81	stupefare (<i>to amaze</i>)	0.82	sgomento (<i>dismay</i>)	0.74
	sgomento (<i>dismayed</i>)	0.73	sconcertare (<i>to disconcert</i>)	0.81	trepidazione (<i>trepidation</i>)	0.74
	allibito (<i>shocked</i>)	0.73	rimanere (<i>to remain</i>)	0.79	turbamento (<i>turmoil</i>)	0.74
	preoccupato (<i>worried</i>)	0.72	indignare (<i>to makeindignant</i>)	0.74	commozione (<i>deep feeling</i>)	0.74
	sconvolto (<i>upset</i>)	0.72	guardare (<i>to look</i>)	0.73	presentimento (<i>presentiment</i>)	0.73
DISGUST	immondo (<i>dirty</i>)	0.6	scandalizzare (<i>to shock</i>)	0.63	fetore (<i>stink</i>)	0.84
	malsano (<i>unhealthy</i>)	0.58	indignare (<i>to makeindignant</i>)	0.53	escremento (<i>excrement</i>)	0.83
	insopportabile (<i>intolerable</i>)	0.58	disapprovare (<i>to disapprove</i>)	0.5	putrefazione (<i>rot</i>)	0.82
	orribile (<i>horrible</i>)	0.56	criticare (<i>to criticize</i>)	0.49	carogna (<i>lowlife</i>)	0.74
	indegno (<i>shameful</i>)	0.52	biasimare (<i>to blame</i>)	0.49	miasma (<i>miasma</i>)	0.74
FEAR	impotente (<i>helpless</i>)	0.6	stupefare (<i>to amaze</i>)	0.7	disorientamento (<i>disorientation</i>)	0.82
	inquieto (<i>restless</i>)	0.57	scioccare (<i>to shock</i>)	0.68	angoscia (<i>anguish</i>)	0.81
	infelice (<i>unhappy</i>)	0.55	sbalordire (<i>to astonish</i>)	0.68	turbamento (<i>turmoil</i>)	0.79
	diffidente (<i>suspicious</i>)	0.53	sconcertare (<i>to disconcert</i>)	0.66	prostrazione (<i>obeisance</i>)	0.79
	spaesato (<i>disoriented</i>)	0.53	disorientare (<i>to disorient</i>)	0.65	inquietudine (<i>apprehension</i>)	0.78
SADNESS	triste (<i>sad</i>)	0.8	deludere (<i>to betray</i>)	0.78	tristezza (<i>sadness</i>)	0.91
	tetro (<i>gloomy</i>)	0.65	amareggiare (<i>to embitter</i>)	0.75	sconforto (<i>discouragement</i>)	0.88
	sconsolato (<i>surrowful</i>)	0.62	angosciare (<i>to anguish</i>)	0.72	disperazione (<i>desperation</i>)	0.88
	pessimistico (<i>pessimistic</i>)	0.61	frustrare (<i>to frustrate</i>)	0.71	angoscia (<i>anguish</i>)	0.88
	angosciato (<i>anguished</i>)	0.59	sfiduciare (<i>to discourage</i>)	0.71	inquietudine (<i>apprehension</i>)	0.87
TRUST	disinteressato (<i>disinterested</i>)	0.65	domandare (<i>to ask</i>)	0.64	serietà (<i>seriousness</i>)	0.91
	rispettoso (<i>respectful</i>)	0.65	dubitare (<i>to doubt</i>)	0.59	prudenza (<i>caution</i>)	0.9
	laborioso (<i>hard-working</i>)	0.64	meravigliare (<i>to amaze</i>)	0.58	mitezza (<i>mildness</i>)	0.89
	disciplinato (<i>disciplined</i>)	0.63	rammentare (<i>to remind</i>)	0.56	costanza (<i>tenacity</i>)	0.89
	zelante (<i>zealous</i>)	0.62	supporre (<i>to suppose</i>)	0.56	abnegazione (<i>abnegation</i>)	0.88
ANTICIPATION	inquieto (<i>agitated</i>)	0.7	sforzare (<i>to force</i>)	0.56	oracolo (<i>oracle</i>)	0.77
	ansioso (<i>anxious</i>)	0.58	confortare (<i>to comfort</i>)	0.56	premonizione (<i>premonition</i>)	0.74
	desideroso (<i>desirous</i>)	0.56	degnare (<i>to deign</i>)	0.55	preveggenza (<i>presage</i>)	0.73
	entusiasta (<i>enthusiastic</i>)	0.56	distogliere (<i>to deflect</i>)	0.55	auspicio (<i>auspice</i>)	0.72
	dubbioso (<i>uncertain</i>)	0.55	appagare (<i>to satiate</i>)	0.54	arcano (<i>aracane</i>)	0.71