

A Deep Learning Semantic Approach to Emotion Recognition Using the IBM Watson Bluemix Alchemy Language

Giulio Biondi¹, Valentina Franzoni²(✉), and Valentina Poggioni¹

¹ Department of Mathematics and Computer Science, University of Perugia,
Via Vanvitelli 1, Perugia, Italy

giulio.biondi@studenti.unipg.it, valentina.poggioni@unipg.it

² Department of Computer, Control and Management Engineering,
Sapienza University of Rome, Via Ariosto, 25, Rome, Italy
franzoni@dis.uniroma1.it

Abstract. Sentiment analysis and emotion recognition are emerging research fields of research that aim to build intelligent systems able to recognize and interpret human emotions. Due to the applicability of these systems to almost all kinds of markets, also the interest of companies and industries is grown in an exponential way in the last years and a lot of frameworks for programming these systems are introduced. IBM Watson is one of the most famous and used: it offers, among others, a lot of services for Natural Language Processing. In spite of broad-scale multi-language services, most of functions are not available in a lot of “secondary” languages (like Italian). The main objective of this work is to demonstrate the feasibility of a translation-based approach to emotion recognition in texts written in “secondary” languages. We present a prototypical system using IBM Watson to extract emotions from Italian text by means of Bluemix Alchemy Language. Some preliminary results are shown and discussed in order to stress pro and cons of the approach.

Keywords: Affective computing · Emotion recognition · Sentiment analysis · SEMO · Semantics · Deep learning · Machine learning · Artificial intelligence · Computer science · IBM · Watson · Alchemy language · Language translation

1 Introduction

In recent years, Affective Computing become a trending topic in Artificial Intelligence focusing, among others, on detection and extraction of affective information from different sources e.g., images, music, brain scans, text. Particularly studied is the analysis of text, written either in structured/semi-structured form or in natural language [27–30, 33]. Using Natural Language Processing (NLP) techniques, automated agents can gain the ability to process and analyze text at different levels of abstraction, exploiting the speed and computational power of modern systems.

Among various branches of Affective Computing, it is possible to identify, Sentiment Analysis (SA) and Emotion Recognition (ER):

- Sentiment Analysis studies the polarity of a text, which can be positive, negative or neutral.
- Emotion Recognition extracts and recognizes the emotional content, according to models of emotions accepted as State of the Art, e.g. the Ekman, Plutchik, and Lovheim models.

Several other models are present in literature and others will come, because the hype initiated at the end of last century didn't still finish its effect on the scientific research [31]. Furthermore, the technical advent of new and affordable tools to measure aspects linked to emotional reactions lead to an interdisciplinary worldwide interest to Affective Computing and Emotion Recognition, so that both private and public research groups are studying emotions in different points of view. For our point of view focusing on textual information, different approaches can be followed to analyze text e.g., knowledge-based, statistical or hybrid approaches. Nowadays, companies already invest a lot of resources to fructify information on the sentiment expressed towards a target, e.g. political subjects, brands, a product with its associated reviews; the efficiency of a marketing campaign depends largely on the ability to detect ongoing trends and to foresee future customers' needs.

1.1 Affective computing for marketing applications

Recommender Systems (RS) can use affective computing to implement and augment knowledge about the sentiment/opinion towards a product or a brand, helping sellers to automatically tailor their recommendations to users with a user-centered approach for a better personalization. Along with commercial applications, researchers are applying Affective Computing also to spam detection [23], and to build human-computer interaction systems based on natural language, both spoken [5,32] and written [13,14,18–21,26].

Another application of Emotion Recognition (ER) that is already on the market, is used to detect in real time the interest, involvement and emotions of a customer through video-based techniques. A simple classification of visual features captured with a camera, recognized and tracked while a user is looking, hearing or trying a product, gives sellers immediate feedback about their products and the connected customers' feelings.

1.2 Novel Applications

Novels applications include to use knowledge of emotions as prosthesis knowledge for tutoring, or auto-evaluation of teachers and leaders, who can analyze their facial or written or spoken expressions in order to be aware of emotions and manage them as possible, with the aim of improving their social skills and avoid human misunderstanding of tone. For example the SEMO model of Semantic Emotion Recognition [1,13] allow to detect emotions from text and thus to eventually correct it before sending it via email or posting on social networks [22].

1.3 A Translation-Based Approach

In this work, a translation-based approach to emotion recognition in texts written in “secondary” languages is presented, and the feasibility of such approach is discussed. For this purpose, we present a prototypical system using IBM Watson to extract emotions from Italian text by means of Bluemix Alchemy Language. Some preliminary results are shown and discussed in order to stress pros and cons of the approach. In the first section of this paper an overview of the software tools and platforms that perform our approach is presented. A detailed breakdown of the IBM Alchemy Language API is given, describing each function related to a complete process of emotion recognition and extraction. The second section shows an example of usage of the Sentiment and Emotion APIs, in conjunction with the Language translator module, used here to translate terms from Italian to English. A final discussion concludes the paper.

2 Platform and Tools

2.1 IBM Watson

IBM Watson is an AI system capable of answering questions posed in natural language. Developed within the *DEEPQA* IBM project, it improves the capability of artificial intelligences in understanding questions posed in natural language, providing relevant answers. Watson was firstly shown in public in 2011, when it competed against two human champions in the TV quiz *Jeopardy!*. The stress was put on the ability of the AI to understand questions and provide exact answers with speed comparable to humans: Watson, in fact, won the first prize in the game.

IBM engineers then focused on using Watson as a clinical decision support system, in partnership with the Memorial Sloan-Kettering Cancer Center, New York. The goal was to enable medical personnel to supply the system with a patient’s clinical situation, including symptoms and risk factors, and get a list of personalized recommendations, ranked by a confidence index. Tests shown that for particular applications (e.g. lung cancer diagnosis) Watson could dramatically improve the time needed for diagnosis and give relevant suggestions. More generally, the key feature of Watson in the medical field is the ability to process the huge amount of knowledge produced in medical literature on a daily basis, at a speed impossible for humans. Diagnoses can then be performed taking into account the whole literature, possibly leading to improved results, which will then be analyzed and used by human personnel.

After the initial presentation of Watson to the public, a variety of Artificial Intelligence applications have been introduced in the service, using machine learning techniques for rapid exploration of an enormous source of linked data [6, 7, 15]. Knowledge in several fields was enhanced, ranging from NLP to personality insights for annotating user profiles basing on the content of their messages, emails and tweets. In a future Internet-of-Things (IoT) vision, General Motors is planning to use Watson to improve location-aware services in their cars, and

the Imperial College in London is focusing on crime prediction using Watson for pattern analysis. We are showing here how Watson can be used also for emotion recognition, in a simple application on text data [8,17].

2.2 IBM Bluemix

IBM Watson tools are available on Bluemix, the IBM cloud computing PaaS (Platform-as-a-Service). Bluemix serves both as an infrastructure, offering the possibility to rent computing power, network appliances and storage space, and as a software platform, with runtime applications developed in different languages and technologies. Bluemix offers access to various services too; a list of all the available services can be found on the Bluemix website (an IBM account is needed).

2.3 Watson Developer Cloud

IBM *Watson Developer Cloud* offers services for helping users to build cognitive applications; it is the access point for developers to the Watson API, and the basis for the service of Affective Computing. A plethora of services are available on the platform, covering five main areas:

1. Language, recognizing 97 spoken languages;
2. Speech, with NLP functions;
3. Vision, with image recognition;
4. Data Insights for data analysis;
5. Embodied Cognition, including sentiment analysis and alchemy language on linked data.

It is possible to combine functions and libraries belonging to different areas, in order to build complex multimedia applications.

2.4 AlchemyLanguage

AlchemyLanguage is a set of services to build complex applications for textual NLP, with a high level of abstraction used in several applications of emotion recognition (i.e. see [16]). This service is part of the Language area in the Watson Developer Cloud. *AlchemyLanguage* is composed of two main modules:

1. *AlchemyLanguage functions*: a set of text analysis functions to extract semantic information from input text terms;
2. *AlchemyData*: a database indexing about 300 thousand news and blog articles everyday, in English. Users can create custom NLP queries that will be performed directly on the cloud database, without the need of heavy calculation on local client machines. Results will include the indexed articles, matching the query constraints and enriched with NLP information.

2.5 Sentiment Analysis and Emotion Recognition

Entity Recognition. The *Entities* function extracts entities from text. The user inputs a text, HTML document or URL pointing to a web page, and Watson returns a list of entities, using complex statistical algorithms and NLP technology in the form of a combination of multilingual support, linked data, context-sensitive entity disambiguation, comprehensive type support and quotations extraction. The API supports custom models that can be trained using Watson Developer Studio. Results can include additional content, e.g.:

- linked data (active by default);
- entity-level sentiment;
- emotions associated to detected entities.

Supported languages for the function are: English, French, German, Italian, Portuguese, Russian, Spanish, and Swedish.

Sentiment on Document. The *Sentiment* function analyses the overall sentiment in a given text. The user inputs some text content, HTML documents or URL pointing to the information to analyze, and the system returns the associated sentiment, in the form of a triple:

1. a flag *Mixed*, with *1* meaning that the sentiment is not polarized (i.e. both positive and negative), and *0* meaning that the sentiment is polarized;
2. a weight in the $[-1; 1]$ range representing the sentiment, with *0* meaning Neutral;
3. a polarity value among: *positive*, *negative* or *neutral*.

Supported languages are: English, French, Italian, German, Portuguese, Russian, and Spanish.

Sentiment on Specific Targets. The *Targeted Sentiment* function works like the previous, but at a different level of abstraction, associating a sentiment to user-specified targets, e.g. a product or a brand. Furthermore, other items can be analyzed as part of other calls, e.g. quotations, entities (*Entities* call), and keywords (*Keywords* call). Thanks to the strong NLP abilities of Watson, target options include also relational and directional functions (*Relations* call). Supported languages for these functions are: Arabic, English, French, German, Italian, Portuguese, Russian, and Spanish.

3 Experimental Use Case: Assigning Emotional Tags to Book Blurbs in Italian Language

In a previous work, the authors of this work tested if it is possible to classify books using emotional by means of classifying terms extracted not from the whole book, but from its presentation blurb (usually in the fourth page of the edition). Experiments has been carried on books written or translated in Italian language, and the final emotions has been compared to the emotional tags given by users of the

social network Zazie. Zazie is an Italian product giving to users a place where to review the books they read, assigning them reviews and emotional tags to discover, from collective information, new books to choose, based on their capacity to raise emotions (e.g. happy, sad) or to be read while in particular moods (e.g. relaxed on the coach, distracted on the beach). In [12, 14] the authors tested the possibility of automatic classification of books associating to the words of blurbs to emotional tags extracted from WordNet Affect, which is part of the WordNet ontology. In this work, we implement a simple experimentation of what the IBM Alchemy API can do for the same task, using a simple example from the Zazie data set. In [9, 11] the Wikipedia semantic network is navigated by using a trace based approach [25] similar to ACO algorithms [2–4, 24] and several web-based proximity measures [10], it is worth investigating in future work if the same techniques and proximity measures can be used in the framework of emotion classification.

3.1 Experiment Design

As shown in the previous section, Alchemy APIs provide the user with methods to assign a sentiment/emotion tag. While sentiment can be extracted in a number of languages, including Italian, emotions can be instead recognized only in English. A translation from Italian to English will be therefore exploited using the *Language Translator* module as a preliminary action, then the result will be used with the *Emotion Analysis* module. As an example, we will carry out the Italian text of the blurb from the book *Il turno di notte lo fanno le stelle* by the author Erri De Luca:

Matthew ha un cuore nuovo. Un cuore di donna. Ed è con una donna, sua compagna di malattia e guarigione in ospedale, che andrà a riprendersi la vita in cima a una montagna, scalandola. Uno straordinario racconto-sceneggiatura, dove ogni battito del cuore ha un suono mai udito prima.

3.2 NLP Preprocessing Phase

The *Relations* function extracts triplets (Subject, Action, Object) expressing relations. It is one of the most advanced and interesting tools available in the AlchemyLanguage API, and well demonstrates the power of Watson Data Mining AI in the field of NLP. The user inputs some text content, HTML documents or URLs pointing to the information to analyze, and the system returns a list of sentences along with their elements, performing the following steps, for which we provide a description including a tag “charged” if they are due to an additional payment in the IBM system:

1. *Sentence Parsing*: parses sentences into subject, action, object and location.
2. *Keyword extraction* (charged): extracts keywords information
3. *Entity extraction* (charged): extracts entities information.
4. *Sentiment Analysis* (charged): calculates sentiment at a relational level for subjects, objects and locations, and directional level for actions.
5. *Verb Normalization*: lemmatizes verbs, identifies the tense and detects an eventual negation.

By default, linked data content links are included in the response; users can optionally hide them.

Supported languages are English and Spanish. Italian is not supported, thus our application will require a preliminary translation.

3.3 Results on the Sentiment Analysis Task

In this section a Sentiment tag is assigned to the blurb, among *positive*, *negative* and *neutral*.

The following script queries the *Sentiment Analysis* API and prints the result, in the form of a JSON document, on screen.

Algorithm 1. Sentiment Analysis example code

```

1 import json
2 from watson_developer_cloud import AlchemyLanguageV1
3
4 alchemy_language = AlchemyLanguageV1(api_key=API_KEY)
5 text="Matthew_ha_un_cuore_nuovo._Un_cuore_di_donna._Ed_
   \ 'e_con_una_donna,_sua_compagna_di_malattia_e_
   guarigione_in_ospedale,_che_andr\'a_a_riprendersi_
   la_vita_in_cima_a_una_montagna,_scalandola._Uno_
   straordinario_racconto-sceneggiatura,_dove_ogni_
   battito_del_cuore_ha_un_suono_mai_udito_prima."
6 result = alchemy_language.sentiment(text=text,
   language=italian)
7 print(json.dumps(result,indent=2))

```

The following JSON-style document shows the result of the call.

```

1 {
2     "docSentiment":{
3         "mixed" : "0",
4         "score" : "0.467522",
5         "type" : "positive"
6     },
7     "usage": "By accessing AlchemyAPI or using information
8     generated by AlchemyAPI, you are agreeing to be bound by the
9     AlchemyAPI Terms of Use:
10    http://www.alchemyapi.com/company/terms.html",
11    "status": "OK",
12    "totalTransactions": "1"
13 }

```

Listing 1: JSON example

Results include some control information which is common to each API call, i.e.:

- *usage* field, containing a link to the terms of usage of AlchemyAPI;
- *status* field, indicating whether the request was successful or not;
- *totalTransactions* field, indicating how many API calls were necessary to get all the required information. May be higher than one, for example in combined calls.

The sentiment information is embedded in the *docSentiment* object, which is composed of three elements:

- A “Mixed” flag, with 1 meaning that the sentiment is both positive and negative, and 0 meaning that the sentiment is polarized
- A weight in the [-1; 1] range representing the sentiment strength, with 0 meaning *Neutral*
- A Sentiment tag that can be Positive, Negative or Neutral.

In our case, the API reported a Positive polarity for the text, with a score of 0.467522 and no mixed sentiment.

Note: the *mixed* attribute is not included in the result when its value is 0. We decided to include it for clarity purposes.

3.4 Translation and Emotion Analysis

In this section we will show the second step of our example, translating the blurb to the English language and calculating emotional scores.

The following script first queries the *Language Translator* service, prints the result on screen and uses it as input to the *Emotion Analysis* API.

Algorithm 2. Translator and Emotion Analysis example code

```

1 import json
2 from watson_developer_cloud import LanguageTranslatorV2 as
   LanguageTranslator, AlchemyLanguageV1
3 language_translator = LanguageTranslator(
4     username='LT_userName',
5     password='LT_password')
6 alchemy_language = AlchemyLanguageV1(api_key=API_key)
7 text=XXX
8 translation = json.dumps(language_translator.translate(text=text,
9     source='it', target='en'), indent=2, ensure_ascii=False)
9 print(translation)
10 result = json.dumps(alchemy_language.emotion(text=translation),
11     indent=2)
11 print(result)

```

The translation performed by the Language Translator service:

Matthew has a new heart. A woman's heart. And it's with a woman, her roommate disease and healing in the hospital, which will recover their lives at the top of a mountain, scalandola. One amazing racconto-sceneggiatura, where each pulse of the heart has a sound never heard before.

The JSON-style document in listing 2 shows the result of the call for Emotion Analysis. The emotion information is embedded in the *docEmotions* object, which contains five emotion-score pairs, as described in Chap. 2, section *Emotion Analysis*. In our case, the API reports *joy* as the prevalent emotion in the text, followed in turn by *fear*.

4 Results Discussion and Conclusion

In this work an overview of the *AlchemyLanguage* API was given and an example use case on the Zazie data set was shown. The main goal of this work was to check if IBM Watson Developer Cloud Services can be used for research purposes in the area of Emotion Recognition and Sentiment Analysis, both as ways to discover new information and ease processes and as benchmarks to compare against laboratory results. The answer is positive, although with some caveats:

```

1  {
2  "docEmotions":
3  {
4    "fear": "0.213389",
5    "joy": "0.301097",
6    "sadness": "0.142693",
7    "anger": "0.200515",
8    "disgust": "0.206219"
9  }
10 "totalTransactions": "1",
11 "language": "english",
12 "usage": "By accessing AlchemyAPI or using information generated
13    by AlchemyAPI, you are agreeing to be bound by the AlchemyAPI
14    Terms of Use:
15    http://www.alchemyapi.com/company/terms.html",
16 "status": "OK"
17 }
```

Listing 2: Emotion analysis results example

- First, even if IBM allows to launch the computation on their dedicated cloud platform, it could prove expensive in terms of time: there is a free daily/-monthly quota for AlchemyLanguage, depending on the different methods for

Language Translator; other services, which we only mentioned, have their own subscription plans. Those amounts are enough for small testing, but working on Big Data may require planning to optimize time: it was calculated that to perform *Sentiment Analysis* on all the blurbs included in the Zazie dataset without incurring in any charge approximately 10 days are needed, with the same time required for *Emotion Analysis*.

- Building and using custom models for specific needs requires a consistent effort in terms of money. The free *one time* quota is 250.000 characters.
- The quality of translation still needs to be improved. In the previously shown example, the system was not able to translate some words, which were left in the original language. This could lead to less precise results, especially in the case when the words

Nevertheless, IBM *AlchemyLanguage*, even in the limited *free* version, represents a good and easy way to access tools at the state of art of Natural Language Processing. It would be interesting to analyze all the services on Watson Developer Cloud and their applications, to test not only text-related but Image Recognition and Speech Recognition ones, to perform Affective Computing and Emotion Recognition from images and audio directly on the IBM cloud.

References

1. Anusha, V., Sandhya, B.: A learning based emotion classifier with semantic text processing. In: El-Alfy, E.-S.M., Thampi, S.M., Takagi, H., Piramuthu, S., Hanne, T. (eds.) *Advances in Intelligent Informatics*. AISC, vol. 320, pp. 371–382. Springer, Cham (2015). doi:[10.1007/978-3-319-11218-3_34](https://doi.org/10.1007/978-3-319-11218-3_34)
2. Baiocchi, M., Milani, A., Poggioni, V., Rossi, F.: An ACO approach to planning. In: Cotta, C., Cowling, P. (eds.) *EvoCOP 2009*. LNCS, vol. 5482, pp. 73–84. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-01009-5_7](https://doi.org/10.1007/978-3-642-01009-5_7)
3. Baiocchi, M., Milani, A., Poggioni, V., Rossi, F.: Ant search strategies for planning optimization. In: *ICAPS 2009 Proceedings of the 19th International Conference on Automated Planning and Scheduling*, pp. 334–337 (2009)
4. Baiocchi, M., Milani, A., Poggioni, V., Rossi, F.: Optimal planning with ACO. In: Serra, R., Cucchiara, R. (eds.) *AI*IA 2009*. LNCS, vol. 5883, pp. 212–221. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-10291-2_22](https://doi.org/10.1007/978-3-642-10291-2_22)
5. Bhaskar, J., Sruthi, K., Nedungadi, P.: Hybrid approach for emotion classification of audio conversation based on text and speech mining. *Procedia Comput. Sci.* **46**, 635–643 (2015)
6. Chiancone, A., Franzoni, V., Milani, A.: A multistrain bacterial diffusion model for link prediction. *Int. J. Pattern Recogn. Artif. Intell.* **31**(11), 157–172 (2017). World Scientific
7. Chiancone, A., Milani, A., Poggioni, V., Pallottelli, S., Madotto, A., Franzoni, V.: A multistrain bacterial model for link prediction. In: *Proceedings of the International Conference on Natural Computation*, pp. 1075–1079. IEEE Press (2016). doi:[10.1109/ICNC.2015.7378141](https://doi.org/10.1109/ICNC.2015.7378141)
8. Ferrucci, D.A.: Introduction to this is watson. *IBM J. Res. Dev.* **56**(34), 1 (2012)

9. Franzoni, V., Mencacci, M., Mengoni, P., Milani, A.: Semantic heuristic search in collaborative networks: measures and contexts. In: Proceedings 2014 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology, WI/IAT 2014, vol. 1, pp. 187–217. IEEE Press (2014). doi:[10.1109/WI-IAT.2014.27](https://doi.org/10.1109/WI-IAT.2014.27)
10. Franzoni, V., Milani, A.: Pming distance: a collaborative semantic proximity measure. In: Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology, IAT 2012, vol. 2, pp. 442–449. IEEE Press (2012). doi:[10.1109/WI-IAT.2012.226](https://doi.org/10.1109/WI-IAT.2012.226)
11. Franzoni, V., Milani, A.: A pheromone-like model for semantic context extraction from collaborative networks. In: Proceedings IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology, WI-IAT 2015, 2016-January, pp. 540–547, IEEE Press (2016)
12. Franzoni, V., Poggioni, V., Zollo, F.: Can we infer book classification by blurbs. CEUR Workshop Proceedings, vol. 1127, pp. 16–19. CEUR WS (2014)
13. Franzoni, V., Biondi, G., Milani, A., Li, Y.: Web-based semantic similarity for emotion recognition in web objects. CoRR abs/1612.05734 (2016)
14. Franzoni, V., Poggioni, V., Zollo, F.: Automated classification of book blurbs according to the emotional tags of the social network zazie. In: 1st International Workshop on Emotion and Sentiment in Social and Expressive Media, ESSEM 2013, CEUR Workshop Proceedings, pp. 83–94. CEUR WS (2013)
15. Gentili, E., Milani, A., Poggioni, V.: Data summarization model for user action log files. In: Murgante, B., Gervasi, O., Misra, S., Nedjah, N., Rocha, A.M.A.C., Tanar, D., Apduhan, B.O. (eds.) ICCSA 2012. LNCS, vol. 7335, pp. 539–549. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-31137-6_41](https://doi.org/10.1007/978-3-642-31137-6_41)
16. Gupta, R.K., Yang, Y.: Crystalnest at semeval-2017 task 4: Using sarcasm detection for enhancing sentiment classification and quantification. In: SemEval: 11th International Workshop on Semantic Evaluation, Aug 3–4, 2017, Vancouver, Canada (to appear)
17. High, R.: The Era of Cognitive Systems: An Inside Look at IBM Watson and How it Works. IBM Corporation, Redbooks, Armonk (2012)
18. Houjeij, A., Hamieh, L., Mehdi, N., Hajj, H.: A novel approach for emotion classification based on fusion of text and speech. In: 2012 19th International Conference on Telecommunications (ICT), pp. 1–6, April 2012
19. Huang, S.I., Chen, Y.S.: Developing document classifiers for recognizing article readers' affects. In: Proceedings of the 2012 International Conference on Information Management (2012)
20. Liberati, C., Camillo, F.: Subjective business polarization: Sentiment analysis meets predictive modeling. In: Catania, B., et al. (eds.) New Trends in Databases and Information Systems. AISC, vol. 241. Springer, Cham (2014)
21. Lupan, D., Bobocescu-Kesikis, S., Dascalu, M., Trausan-Matu, S., Dessus, P.: Predicting readers' emotional states induced by news articles through latent semantic analysis. In: SMART 2013 International Conference on Social Media in Academia: Research and Teaching, pp. 79–84. Citeseer (2013)
22. Mancini, L., Milani, A., Poggioni, V., Chiancone, A.: Self regulating mechanisms for network immunization. *AI Commun.* **29**(2), 301–317 (2016)
23. Markines, B., Cattuto, C., Menczer, F.: Social spam detection. In: Proceedings of the 5th International Workshop on Adversarial Information Retrieval on the Web AIRWeb 2009, p. 41 (2009)
24. Milani, A., Poggioni, V.: Planning in reactive environments. *Comput. Intell.* **23**(4), 439–463 (2007)

25. Pallottelli, S., Franzoni, V., Milani, A.: Multi-path traces in semantic graphs for latent knowledge elicitation. In: Proceedings of International Conference on Natural Computation 2016-January, pp. 281–288. IEEE Press (2016). doi:[10.1109/ICNC.2015.7378004](https://doi.org/10.1109/ICNC.2015.7378004)
26. Ren, F., Quan, C.: Linguistic-based emotion analysis and recognition for measuring consumer satisfaction: an application of affective computing. *Inf. Technol. Manage.* **13**(4), 321–332 (2012)
27. Shelke, N.: Approaches of emotion detection from text. *Int. J. Comput. Sci. Inf. Technol. Res.* **2**(2), 123–128 (2014)
28. Shivhare, S.N., Garg, S., Mishra, A.: Emotionfinder: detecting emotion from blogs and textual documents. In: 2015 International Conference on Computing, Communication & Automation (ICCCA), pp. 52–57. IEEE (2015)
29. Shivhare, S.N., Saritha, S.K.: Emotion detection from text documents. *Int. J. Data Min. Knowl. Manage. Process* **4**(6), 51 (2014)
30. Taboada, M., Brooke, J., Tofiloski, M., Voll, K., Stede, M.: Lexicon-based methods for sentiment analysis. *Comput. Linguist.* **37**(2), 267–307 (2011)
31. Vallverdú, J., Trovato, G.: Emotional affordances for humanrobot interaction. *Adapt. Behav.* **24**(5), 320–334 (2016)
32. Vanzo, A., Croce, D., Castellucci, G., Basili, R., Nardi, D.: Spoken language understanding for service robotics in Italian. In: Adorni, G., Cagnoni, S., Gori, M., Maratea, M. (eds.) AI*IA 2016. LNCS, vol. 10037, pp. 477–489. Springer, Cham (2016). doi:[10.1007/978-3-319-49130-1_35](https://doi.org/10.1007/978-3-319-49130-1_35)
33. Wang, H., Xu, H., Liu, L., Song, W., Du, C.: An unsupervised microblog emotion dictionary construction method and its application on sentiment analysis. *J. Inf. Comput. Sci.* **12**, 2729–2739 (2015)