



HAL
open science

Analyzing Emotional States Induced by News Articles with Latent Semantic Analysis

Diana Lupan, Mihai Dascalu, Stefan Trausan-Matu, Philippe Dessus

► **To cite this version:**

Diana Lupan, Mihai Dascalu, Stefan Trausan-Matu, Philippe Dessus. Analyzing Emotional States Induced by News Articles with Latent Semantic Analysis. Proc. 15th Int. Conf. on Artificial Intelligence: Methodology, Systems, Applications (AIMSA 2012) , 2012, Varna, Bulgaria. pp.59-68, 10.1007/978-3-642-33185-5_7. hal-01487441

HAL Id: hal-01487441

<https://hal.science/hal-01487441v1>

Submitted on 12 Mar 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Analyzing Emotional States Induced by News Articles with Latent Semantic Analysis

Diana Lupan¹, Mihai Dascalu^{1,2}, Ștefan Trăușan-Matu¹, and Philippe Dessus²

¹ Politehnica University of Bucharest, Computer Science Department, Romania

² LSE, UPMF Grenoble-2 & IUFM-UJF Grenoble-1, France

diana.lupan@cti.pub.ro,

{mihai.dascalu, stefan.trausan}@cs.pub.ro,

philippe.dessus@upmf-grenoble.fr

Abstract. Emotions are reflected both in verbal and written communication. If in the first case they can be easier to trace due to some specific features (body language, voice tone or inflections), in the second it can be quite tricky to grasp the underlying emotions carried by a written text. Therefore we propose a novel automatic method for analyzing emotions induced by texts, more specifically a reader's most likely emotional state after reading a news article. In other words, our goal is to determine how reading a piece of news affects a person's emotional state and to adjust these values based on his/her current state. From a more technical perspective, our system (*Emo2* – Emotions Monitor) combines a context independent approach (actual evaluation of the news employing specific natural language processing techniques and Latent Semantic Analysis) with the influences of user's present emotional state estimated through his/her specific feedback for building a more accurate image of a person's emotional state.

Keywords: emotional state, Latent Semantic Analysis, automatic evaluation of news articles.

1 Introduction

Emotions are one of the defining elements of human nature, as they are present in daily life and almost in any context. They add value to our interactions in a unique way and their induced effect can change the meaning of an entire message.

Although emotions are observably present in face-to-face or verbal communication, they can also be traced in written messages. If in the first case emotions are mostly expressed using body language and different voice features (for example, tone, rhythm, frequency), in the second one all these features are eliminated, as the communication channel cannot support them. Therefore, it can be difficult to detect emotions in written communication due to the lack of characteristics and our sole elements of analysis are words, grouped into sentences. Despite this fact, analyzing emotions in written texts can provide a better understanding of the underlying message by predicting its most likely impact on the person reading it.

Our approach covers the automatic evaluation of a person's mood and sentiment change when he/she reads an article. Our focus is to analyze short news consisting of

a headline and a brief description that are written with the intention of “provoking” emotions and attracting the reader’s attention.

This paper covers two approaches implemented in the Emo2 (Emotions Monitor) system: *context independent* and *environment-influenced* evaluation of emotional states. The context independent approach implies an evaluation of the article’s content with different natural language processing techniques from the perspective of the intended emotional state induced by the author of the article, while the environment-influenced approach takes into account the user’s current emotional state for adjusting the results, based on the influence of other similar articles to one’s current emotional state.

The remainder of the paper first presents similar related works and a general overview of the architecture describing each module. The fourth section focuses on the context independent evaluation process including details about the word valence and article similarity approaches. The fifth section presents the environment-influenced evaluation method, while the last sections comprise the evaluations, results and conclusions of our survey.

2 Related Work

The UA-ZBSA system [8] aimed at classifying a set of news headlines into six types of emotions: “anger”, “disgust”, “fear”, “joy”, “sadness”, “surprise”. Our approach is similarly comprised of six types of emotions, but we modified them to best fit our dimensions of analysis. Therefore, the included categories include “in control of the situation”, “fear”, “joy”, “sadness”, “excitement” and “boredom” described in detail in the next sections). UA-ZBSA’s emotion classification uses a bag of word approach comprising of frequency and co-occurrence counts of concepts collected from the World Wide Web. It applies the idea that “words which tend to co-occur across many documents with a given emotion are highly probable to express that emotion”. Moreover, it considers that adjectives with the same polarity tend to appear together.

On the other hand, UPAR7’s design [9] is based on the assumption that all the words from an article’s title potentially carry emotions. The objective of the system was to identify the expression which carries the main topic of the title as this expression has major importance. Moreover, rules for detecting specific emotions were also developed (ex. a negation may be an indicator of a possible surprise). The system consists of the SS-Tagger (a Part-of-Speech tagger) and the Stanford Parser and used WordNet, SentiWordNet and WordNet-Affect as lexical resources. It also counts lexical elements that are later on classified as good indicators of surprise: negations, modal auxiliaries and question marks.

3 General Overview of the Functional Architecture

The designed system consists of several modules, grouped into processing layers that are presented in Fig. 1:

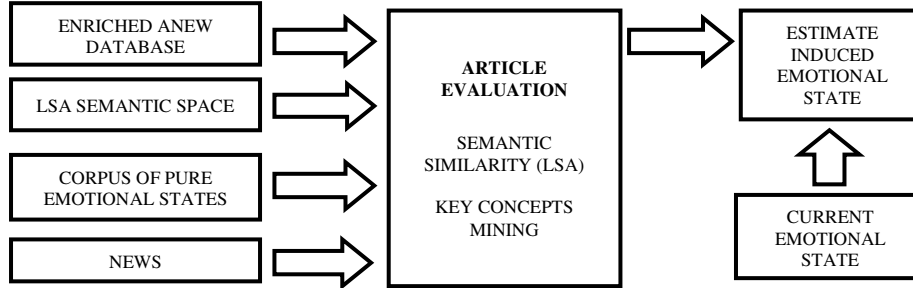


Fig. 1. System's Architecture

Firstly, regarding the overall processing algorithm, we must determine a person's initial or current emotional state. Therefore, our application tries to determine as accurately as possible the user's current emotional state by analyzing the feedback given by him on a set of 5 news articles, with regards to his/her induced current state. Although these initial news might alter a person's current emotional state, the results are later on used for adapting and personalizing future automatic results provided by the system through regressions in order to best reflect the actual impact of each new article on the user's current state.

From a higher level perspective, our analysis is divided into two separate steps: the initial, objective assessment, performed in the "Core evaluation" module, and the subjective, adaptive component of personalizing the results by considering the current emotional state of the user (integrated within the "Estimate Induced Emotional State" component).

The proposed evaluation method uses multiple inputs: the enriched ANEW database, a trained LSA semantic vector space, a corpus of articles expressing pure states and the news currently analyzed. The enriched ANEW database [1] consists of concepts (words) that have three values associated expressing its emotional state (Happy/ Unhappy, Calm/ Excited, Controlled/ In-Control). The Happy/ Unhappy dimension is quite straightforward as it expresses a state of happiness or sadness. Calm/Excited refers to the intention of the user to keep reading about the information presented in the current news article. The last dimension underlines the feeling of being in control of the situation just read versus being controlled by the events (or persons) described in the present news. The values for each dimension are in the [1, 9] range and the initial database of approximately 1,000 words was enriched using an algorithm that consists of 2 stages that are presented in detail in the following section.

On the other hand, the evaluation process can be divided in two approaches: **semantic similarity** evaluation between texts by means of LSA and **mining key concepts**. The later consists of extracting key concepts from both the body and the title of an article and combining their "emotional values" in order to determine the overall induced emotions. Therefore, a set of reference documents (initially the corpus consisted of 10 "pure" documents per emotional state) is used to obtain the semantic dimension of the analysis, while key-concepts mining offers the lexical dimension of our analysis. Key-concept mining is applied separately on the title and on the body of

news because we empirically considered that the title contains more valuable information than the rest. The title is usually designed to be short and to attract the reader's attention; moreover, we can consider that the title expresses from a human perspective the essence and the synthetic message of the body. Thus, due to their intrinsic value, the key-concepts belonging to the title have a greater contribution in the formula than the ones extracted from the body.

In the process of evaluating the emotional state by means of Latent Semantic Analysis, the result is determined by using the average value of all documents found similar with the content of the currently assessed article, weighted by the similarity between them. By using the similarity of concepts that induce a certain state, we can imply that documents with similar structure in terms of coverage of key concepts inflict a similar emotional state. The previous statement is not generally valid, but under the consideration of comparing news articles from a specific domain, within a short timeframe, our results from this approach proved to be quite accurate and to level the spikes of different dimensions of certain news articles within a predefined timeframe. In essence, news articles covering the same topic in a short timeframe rarely change their messages in terms of induced emotions.

The values obtained from the two perspectives are combined so the final result approximates, in the best way possible, the emotional state that is most likely to be induced when a person reads the article. The "Current emotional state" module integrates the user's input in our analysis (a questionnaire and 5 news articles) by estimating his present state and the fluctuations observed after reading each article. This set of articles contains texts extracted from the latest news feed into our application so their influence on the assessment of other similar articles can be considered relevant and significant. In our experiments, articles consisting of the title and the content (typically 2 or 3 phrases) were extracted from CNN.com's RSS (Really Simple Syndication) feeds.

4 Context-Independent Evaluation

The context-independent approach is based on two different stages: the first compares the concepts found in the analyzed article with terms stored in our enriched ANEW database (Word valence approach), while the second determines the similarities between the article and a set of predefined documents that express a dominant emotional state (Article similarity). In the end, the two approaches are combined during the article evaluation step for obtaining the corresponding valences on all 3 dimensions as a final score for each analyzed news article.

4.1 Word Valence Approach

Each term is associated with three connotations that express what kind of feelings it arouses when someone reads it. These valences expressed on the 3 dimensions (Happy/ Sad, Calm/ Excited, Controlled/ In-Control) are into a local database, formed through an enrichment process that is described in detail in the following paragraphs.

Initially, the number of stored terms was around 1,000 and the associated values were obtained from the ANEW database (Affective Norms for English Words, [1]) that was developed to provide a set of normative emotional ratings for a large number of words in English. This initial database was afterwards increased by using WordNet [3] synsets and Latent Semantic Analysis – LSA [4, 7] that enabled the determination of similarities between a “known” concept and a “new” one. By combining metrics of similarity from a lexicalized ontology with cosine similarity between concepts of the semantic vector space [5], our results were improved and we were able to better weigh the valence of words not present in the initial ANEW database.

LSA was trained on the Reuters corpus of news that provided adequate materials, relevant from the perspective of topics with regards to currently assessed articles and general enough to grasp underlying connections between concepts. The main reasons of using LSA are that its behavior has been proven to be close to that of human beings [2] and that it can be successfully used to enrich data obtained from experiments with human subjects and to approximate how reading a news article affects someone’s emotional state [6].

The actual enrichment algorithm consists of two stages. During the first stage we select a word and determine its synonyms using synsets from WordNet, whereas in the second stage similar concepts, from similar contexts, are determined through LSA. By considering the cosine similarity as the semantic metric for comparing two concepts, we can introduce new terms into our enriched database based on “known” concepts (already existent in the database), their corresponding values and the similarity with the term to be added. Moreover, only “known” concepts that are similar over a certain threshold are considered synonyms and are taken into account in computations. Therefore, based on the synonyms identified through WordNet of the concepts already present in the ANEW database and on the semantic similarity with k most nearest concepts, we determine a new word’s valence and introduce it in our enriched database. After running multiple iterations with increasing k values, in our final experiments we chose k to be the *three* most dominant and similar concepts by means of cosine similarity to the currently assessed one.

4.2 Article Similarity

As previously mentioned, this approach consists of determining similarities between the content of a news article with a corpus of documents that express a pure emotional state. Although the concept of “pureness” is subject to personal evaluation, these selected documents express a dominant dimension, identified and agreed upon by several human evaluators. The initial corpus used within our experiments comprises of 10 general interest news articles per state and LSA was used to compute the similarity between the currently evaluated article and each of the documents within the corpus.

4.3 Article Evaluation

The final result for each article is obtained by combing the values of each of the two previous approaches. Firstly, the concepts from the title are considered to be more valuable, in terms of expressed emotions, than the ones from the body.

The evaluation formula from the word valence approach is the following:

$$value_{key-concepts} = p * TitleValue + (1 - p) * ContentValue \quad (1)$$

where *ContentValue* and *TitleValue* are a linear combination of valences of key-concepts contained within the title and the body of the article.

The value of p was determined experimentally using multiple iterations within the range [0.55, 0.9], with increments of 0.05. The experiments revealed that 0.6 is the value that best combines the message expressed by the title with the one implied by the content and, in this case, the best overall correlation was obtained between the computer-based results and the human evaluators. On the other hand, the content of an article is considered to be similar to a text from the set if the similarity exceeds a threshold of 0.2 (the similarity can take values in the [-1, 1] range). For each emotional state, the arithmetic mean between the values that meet this requirement is computed. Next, the results are linearly normalized in order to be projected on a [1, 9] value range. The final result is computed using the following weighted formula which is designed to give more credit to the value obtained from analyzing the key concepts than to the value determined through article similarity, intrinsically limited by the bag of words approach:

$$finalResult = p * value_{key-concepts} + (1 - p) * value_{setOfDocuments} \quad (2)$$

The best fit for p was experimentally determined by testing decreasing values within the range [0.55, 0.9] and the experiments revealed that $p = 0.65$ leads to the best combination of the values associated with the word valence approach and with the semantic similarity between news and pure emotional state documents.

5 Environment-Influenced Evaluation

The environment-influenced approach comes as an adjustment to the results obtained in the first stage (context-independent evaluation) and can be seen as a feedback loop used for further personalizing the results provided by our system. Our assumption is that the emotional state of a person is differently influenced by a news article depending on his/hers current state and a set of preconceptions (mentality). For example, a person who is initially happy and reads sad news will be less affected and therefore less sad than a person who was initially unhappy. Moreover, the effects of reading the news also depend on specific features of the person reading the article, such as psychological profile, how geographically close he/she is to the events presented in the news.

Taking into account these observations, some additional features which rate these aspects were introduced. Firstly, the current emotional state is assessed by asking the reader to rate 5 news articles which are selected in order to be as similar as possible to the news that will be afterwards automatically analyzed and displayed (Fig. 2); later on, the user must answer a question referring to his current geographical location. The five presented articles are chosen by sorting all the news in the database after their similarities with the main news and picking up the top five most similar ones. Although these initial articles might change the user's current emotional state, their automatic assessment compared to the user's feedback provides valuable information on the current emotional state and offers the basis for applying regressions in order to fine-tune the results for the newly assessed articles.

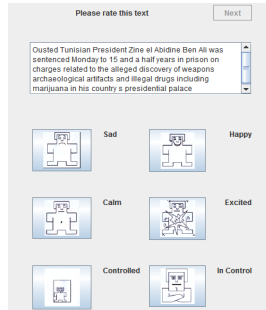


Fig. 2. Graphical interface for the initial rating of articles

The current geographical location of the user is used to adjust the results, following the assumption that events that happen nearby a person tend to affect him/her in a greater proportion. Concepts from the news article are used to determine whether the depicted events are located on the same continent introduced by the user. If the case, valences are augmented relatively to the mean value of each dimension. In order to achieve this functionality, GeoNames (<http://www.geonames.org/>) web services were used for determining the location of events.

6 Results

As mentioned in the previous sections, the designed system was tested on RSS feeds from CNN.com. To test the application, the user is prompted with a window which displays an article and three charts that represent the most likely emotional state induced. Each chart is positioned between two images that suggest a pure emotional state so the user can have an accurate representation of the result. Fig. 3 presents a sample of the application's graphical user interface, including a generated radar graph depicting the 3 dimensions of the analysis:

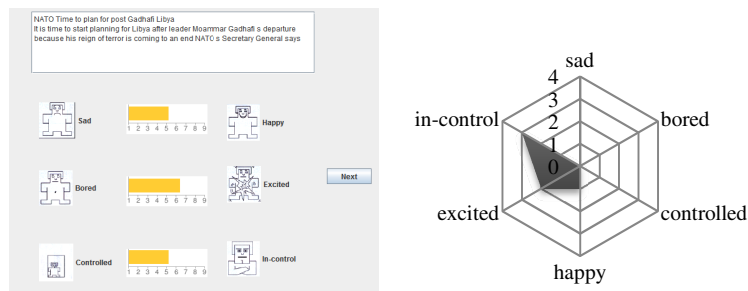


Fig. 3. Graphical interface providing feedback on a specific article and an example of a generated graph

Our validation survey included 10 participants (5 women and 5 men) of age between 20 and 25 with diverse backgrounds (computer science and medical). The survey consisted of 2 stages, each having: 5 initial evaluation news (as previously mentioned, these articles are selected to be as similar as possible to the news to be automatically evaluated by the application) in order to determine the current emotional state, a question referring to the user's geographical position and 10 additional news articles as assessment for obtaining feedback from users regarding the application's results.

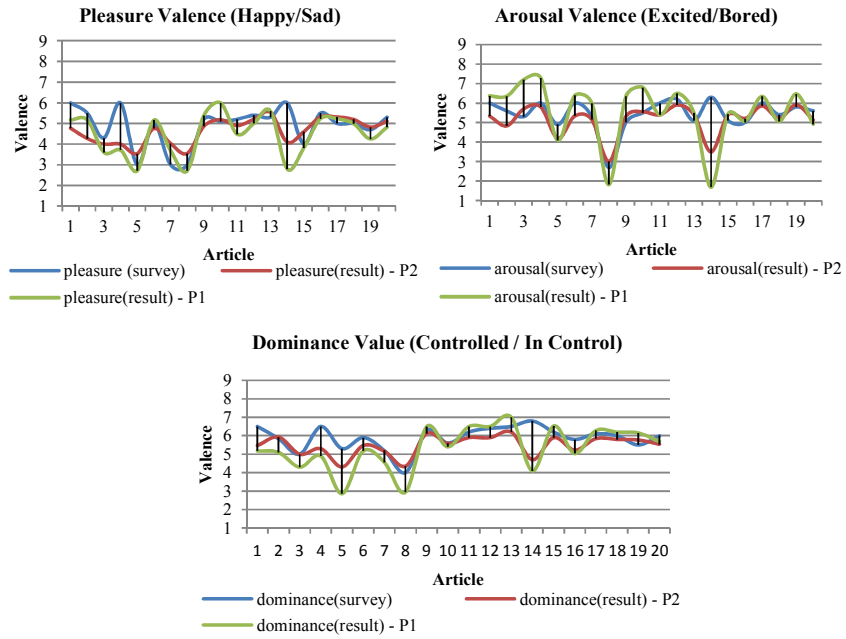


Fig. 4. Pleasure, Arousal and Dominance valences. P1 denotes the context-independent evaluation; P2 denotes the environment-influenced approach.

Therefore, the tests consisted of 20 RSS feeds from all the categories mentioned above. Synthetic results, covering the three dimensions of the analysis (Pleasure Valence, Arousal Valence, Dominance Valence), are presented in Fig. 4.

The correlations between the approaches and the survey, for all dimensions, are presented in Table 1.

Table 1. Correlations between the automatic results and the survey data

	Pleasure Valence	Arousal Valence	Dominance Valence	Average on all 3 dimensions
Survey - P1	0.55	0.48	0.61	0.55
Survey - P2	0.54	0.58	0.56	0.56
P1 - P2	0.86	0.91	0.96	0.91

As it can be observed from the previous table, Pleasure gets equal correlation between the two approaches and the average correlations on all three dimensions are very similar. But, by also taking into consideration the nature of each dimension, we can extrapolate from our experiment that Arousal is a more environment-dependent value in the sense that arousal may vary more than the others depending on the previous mood states, whereas Dominance is more influenced by the actual context and the topics within the article. Additionally, we obtained very high correlations between the two approaches, which was expected due to the fact that P2 refines the initial results.

On the other hand, there are some situations where the second approach obtains significantly inaccurate results in comparison to the survey (that expresses the actual emotions of the user) and also to the other approach. Besides the subjectivity of the humans' evaluation of their own emotional state, a possible explanation for this situation could be the fact that the news from the questionnaire were insufficiently similar to the ones displayed in the results; so, instead of adjusting the values from the first stage to be closer to reality, the algorithm is misguided and modifies them incorrectly, making the final results even worse. A possible solution to this problem could be enriching the news database with more texts, in order for new articles to be evaluated in comparison to other older, but more similar ones.

Additionally, although location provides some valuable input to our analysis, further factors will be addressed in future versions of our system: topics and personal interests highly impacting one's emotional state, gender and similar factors that can be used to differentiate assessment for different categories of users.

7 Conclusions

Our goal was to develop an application that estimates how reading a piece of news affects a person's emotional state and our approach consisted of a context-independent and an environment-influenced analysis of emotions. As technical background, we used several natural language processing techniques, mostly centered on Latent Semantic Analysis. The results were validated through a survey and, through the provided feedback, we can consider that our method is promising and relevant, in an area where subjectivity is the central player. On the other hand, there are some aspects that can be improved; therefore we plan to perform a deeper analysis and understanding of the text and to model the user's emotional state evolution.

As future research, we envision conducting surveys with an increased number of RSS feeds and audience so the obtained feedback would be more accurate. Also, increasing our corpus of "pure" emotional texts and including more psychological factors in our adaptive model could lead to improved results.

Acknowledgment. This work was partially supported by project 264207, ERRIC-Empowering Romanian Research on Intelligent Information Technologies/FP7-REGPOT-2010-1.

References

1. Bradley, M.M., Lang, P.J.: Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings. Gainesville (FL): The Center for Research in Psychophysiology, University of Florida, Tech. Report (1999)
2. Foltz, P.W.: Latent semantic analysis for text-based research. *Behavior Research Methods, Instruments and Computers* 28(2), 197–202 (1996)
3. Fellbaum, C.: *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge (1998)
4. Landauer, T.K., Foltz, P.W., Laham, D.: An introduction to Latent Semantic Analysis. *Discourse Processes* 25(2/3), 259–284 (1998)
5. Manning, C., Schütze, H.: *Foundations of statistical Natural Language Processing*. MIT Press, Cambridge (1999)
6. Wang, L., Wan, Y.: Sentiment Classification of Documents Based on Latent Semantic Analysis. In: Lin, S., Huang, X. (eds.) *CESM 2011, Part II. CCIS*, vol. 176, pp. 356–361. Springer, Heidelberg (2011)
7. Wiemer-Hastings, P.: How latent is Latent Semantic Analysis? In: *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI 1999)*, vol. 2 (1999)
8. Kozareva, Z., Navarro, B., Vazquez, S., Montoyo, A.: UA-ZBSA: A Headline Emotion Classification through Web Information. In: *Proceeding of the 4th International Workshop on Semantic Evaluations, SemEval 2007*, pp. 334–337 (2007)
9. Chaumartin, F.-R.: UPAR7: A knowledge-based system for headline sentiment tagging. In: *Proceeding of the 4th International Workshop on Semantic Evaluations, SemEval 2007*, pp. 422–425 (2007)