Towards Emotion-based Reputation Guessing Learning Agents

Jones Granatyr, Jean Paul Barddal,
Adriano Weihmayer Almeida, Fabrício Enembreck
Graduate Program in Informatics (PPGIa)
Pontifícia Universidade Católica do Paraná (PUCPR)
jones.granatyr@pucpr.edu.br,
{jean.barddal, awalmeida, fabricio}@ppgia.pucpr.br
Curitiba, Brazil

Adaiane Pereira dos Santos Granatyr
Graduate Program in Psychology
Universidade do Contestado (UnC)
adaiane@unc.br
Porto União, Brazil

***Abstract*—Trust and reputation mechanisms are part of the logical protection of intelligent agents, preventing malicious agents from acting egotistically or with the intention to damage others. Several studies in Psychology, Neurology and Anthropology claim that emotions are part of human's decision making process. However, there is a lack of understanding about how affective aspects, such as emotions, influence trust or reputation levels of intelligent agents when they are inserted into an information exchange environment, e.g. an evaluation system. In this paper we propose a reputation model that accounts for emotional bounds given by Ekman's basic emotions and inductive machine learning. Our proposal is evaluated by extracting emotions from texts provided by two online human-fed evaluation systems. Empirical results show significant agent's utility improvements with $p < .05$ when compared to non-emotion-wise proposals, thus, showing the need for future research in this area.**

I. Introduction

Trust and reputation mechanisms consist on observations of agent's behaviors and the assignment of trust values to each one of them, which are later used to aid decision making about whether an agent should or not interact with another. These mechanisms are part of the logical protection of a Multi-agent System (MAS) and are used to prevent malicious agents from acting egotistically or with the intention to spoil others.

Developing intelligent agents based on trust and reputation techniques requires embedding a set of dimensions and characteristics that are essential for these mechanisms to operate. An important dimension is the paradigm type, which expresses the architecture used to build the model, that can be numeric, cognitive or hybrid [1]. Numeric models are based on numerical aggregations of past interactions and employ statistics and probabilities. Conversely, cognitive models are based on beliefs, mental states and mental consequences, thus, related to the BDI (belief, desire and intention) architecture [2]. Finally, hybrid approaches embed both numeric and cognitive traits.

Trust and reputation are often studied in other areas, including Psychology, Neurology and Anthropology [3], [4], [5], since psychological states are claimed to influence human's decision making process. In this context, trust has been divided in cognitive and affective. The first concerns the numeric paradigm and relies on equations to compute trust, while affective trust requires knowledge from emotional processes and is claimed to be more important in decision making processes [6].

So far, there is a lack of understanding about how affective aspects (emotions) influence the level of trust or reputation an agent has when it is inserted into an automated system for information exchange, such as an evaluation system. Although there are models presenting affective characteristics, most of them (i) employ the BDI architecture [7] or (ii) are meta-models constructed upon fixed numeric models [8], [9]. Until this point, trust computation based on emotions is thus unclear [10] and most proposed techniques do not rely on more sophisticated theories, e.g. Appraisal Theories like the OCC model [11] or the Ekman's emotions [12].

In this paper we propose a machine learning-based intelligent agent model that accounts for emotional bounds given by Ekman's basic emotions [12]. We hypothesize that by using this model, an agent will achieve higher correlations when compared to those that do not account for emotional features in an item reputation prediction task. Our proposal is evaluated by extracting emotions from texts provided in two online human-fed evaluation systems. Empirical results show significant correlation improvements with $p < .05$.

This paper is divided as follows. We start by discussing about trust and reputation (Sec. II). Later, we focus on emotions and how to perform their extraction from texts (Sec. III). Following, we show how trust and emotions can participate on decision-making processes (Sec. IV), thus enabling the introduction of our proposal (Sec. V). Since our proposal is generic, it is evaluated with several inductive learning algorithms in two real-world datasets (Sec. VI). Finally, we conclude this paper and state future works (Sec. VII).

II. Trust and Reputation

Trust has been discussed by several disciplines broadening Economy, Philosophy, Sociology and Psychology [13], [14], [15], [16], [17]. More recently, several Computer Science applications, e.g. multi-agent systems, e-commerce, grid computing, recommender and evaluation systems, have demonstrated

978-1-5090-0620-5/16/$31.00

the need to account for these mechanisms, therefore, also drew attention to this field [1], [18].

Trust was initially explored in Sociology and was coined as an attitude related to uncertainty, complexity and inability to predict the future. Authors in [19] define it as an attitude grounded in beliefs about characteristics of other party and elements of a particular situation. Therefore, partners' beliefs are derived from past experience, thus, trust consequences are reflected through the intent to act. Besides this definition, authors in [20] suggested trust as a psychological state comprising the intention to accept vulnerability based on the positive expectations of the forethoughts or behaviors of another.

We refrain from providing an extensive mapping of all possible trust definitions since it has been recently surveyed in [1], yet, to the remainder of this paper, one must differentiate two types of trust: cognitive and affective. Cognitive trust relates to beliefs based on opinions or knowledge about objects. It can be quickly constructed and is composed by a rational content [21], [19]. Essentially, cognitive trust involves conscious decisions about trust partners based on competence, responsibility and dependence [22], [6], [23], [24]. Conversely, affective trust relates to strong emotional content given by the level of care and concern partners exhibit, all based on emotional bonds between individuals [25], [6].

Although both types of trust have shown influence in the anticipation of future events [26], cognitive trust concerns to reasoning mechanisms that enable non-emotion-wise probability-based predictions, therefore it is easier to build and manage. On the other hand, affective trust is grounded on emotions and affective aspects, thus, it is claimed to be more important since it helps to model people's sentiments during decision-making.

Another relevant concept is reputation, which is defined as the expectations held by others based on someone's past behaviors indicated by the members of a community. Similarly, authors in [27] address beliefs about skills and honesty based on recommendations given by others, while [28] argues that reputation helps to manage the complexity of social life since it highlights trustful people who have interest in keeping promises. Reputation acts as a collective memory mechanism that represents an agent's past behavior working as a base for the construction of mutual trust. Reputation thus keeps bad behavior agents away from others, working as a security layer of the system.

For instance, when forming coalitions, agents can benefit from reputation to select the best fit for a specific task raising the utility of the group as a whole [29]. In negotiation and competition scenarios [30], reputation works as a base to the trust that other agent will deliver what is expected from them, or as a defense mechanism that averts possibly bad agents. Reputation can also be used to select information sources [31] or participate in the composition of argumentation models that can trigger standard actions based on agents' reputation levels [32]. For more information about trust and reputation models applied to Multi-Agent Systems, the reader is referred to [1], where the authors discuss about trust and the types of interactions often found in intelligent agents systems.

III. Emotions

Emotions are part of decision making and are considered as instinctive responses when someone faces a situation, thus, they have short duration since they are triggered by certain timely events. Most of people's behaviors are determined by emotions, which generates motives for actions [33]. Emotionless, people would not perceive significance in events. For example, if a purchase of a given product results in dissatisfaction due to its low quality, a temporary anger emotion could be triggered towards the seller.

Aiming at modeling these chains of events, scientists developed Appraisal Theories, which consider emotions as the result of the evaluation of events that are happening in a given time [34]. So far, the most used model in Computer Science is the OCC [11], which assumes that emotions are developed given a set of cognitions and interpretations. This model encompasses 22 emotions, such that half are positive and the other half, negative. Another model is Ekman's basic emotions [12], which is composed of the following emotions: happiness, sadness, fear, surprise, anger, and disgust. Ekman's original study was based on facial appearance, where authors concluded the existence of the latter set of emotions.

A. Emotions and Texts

Besides the increasing effort put on extracting emotions from texts, this task is still labyrinthine. There are two main lines of research regarding this task. The first broads sentiment analysis, which is the extraction of polarity or valence aspects from texts, thus indicating whether a phrase or text is positive, negative or sometimes, neutral. For example, in [35], authors aimed at determining whether a review was positive or negative using supervised machine learning, while the work of [36] performed classification of eBay[1] sellers' reviews.

Conversely, emotion analysis tackles the extraction of fine grained emotions from texts, e.g. happiness, sadness, fear, disgust, anger, and surprise. As an example, authors in [37] aimed at identifying emotions using Appraisal Theories with the OCC model [11], while authors in [38] identify Ekman's emotions on news using latent semantic analysis. Additionally, in [39] authors proposed Synesketch, which also adopts Ekman's basic emotions to detect emotions in short messages. For more information about emotion extraction from texts, the reader is referred to [40] and [41], where the most used computational approaches for this task are surveyed, as well as related work.

Synesketch [39] is composed by several modules, such as Wordnet [42], lexicons, emotional lexicons, abbreviations, colloquialism, and heuristic rules. It estimates emotional weights $w \in [0;1]$ for each possible emotion, which corresponds to their intensity. It also assesses a general weight that indicates the intensity of all identified emotions, as well as the overall

[1] http://ebay.com

phrase valence which is categorized as positive, negative or neutral. The authors found limitations in fear and surprise emotions, while the remaining emotions showed correlations of about 80%. In our experiments we adopted Synesketch since it works with the English language and authors offer it as an open-source API.

Synesketch [39] was built to work on sentence level, so it is not expected to provide good results if a review has more than 140 characters. As evaluation systems allow users to write long texts, it is necessary to separate an entire review in a set of phrases in order to get proper results. To demonstrate this process, let us consider the following review: *"I came here with my boyfriend in January. So convenient when you stay in NYC for a week and don't want to eat out every meal, every day. Anyway, I'm a priority club member so they bumped me up to a room with a nicer view without asking :)"*.

Firstly, the text is split into disjoint phrases and Synesketch is applied separately to each one of them. The latter text review is presented on Table I, where it was divided into three phrases. Emotions and valences were extracted separately from each sentence. A mean is calculated to get the final result about the presence of certain emotions on reviews. For example, for this review the quantity of *fear* is 2 because this emotion is presented in two phrases. On the other hand, *fear weight* is 0.065 because it represents the mean weight of this emotion, i.e. $\frac{0.06+0.07}{2}$. The first phrase is discarded since Synesketch was unable to detect emotions in it.

IV. On Trust, Emotions and Decision-making

In this section we justify the usage of affective aspects in decision making in conjunction with a trust and reputation model. Most mechanisms to compute trust are based on utility functions that do not account for individual characteristics, thus, are claimed to fail since emotions can affect how people perceive and trust one another. Similarly, in [43] authors suggest that emotions are part of decision-making and their integration in an agent's architecture allow more lifelike behaviors. Authors provided empirical evidences that consumers rely on their emotional bounds to service providers as a quality indicator, which fits the definitions of affective trust [44].

Emotions and mood affect the way people perceive things, interpret one's actions, plan, and make decisions [45]. The adoption of psychological architectures in agents are useful to (i) develop cognitive models more similar to humans, (ii) provide the understanding of situations using emotional structures; and (iii) turn the decision making process more efficient and trustworthy [1]. Since emotions are claimed to affect the way people evaluate users or items in evaluation systems, we hypothesize that emotional features aid machine learning-based agents into the task of predicting reputation.

V. Proposed Model

Our proposal predicts the reputation of items using emotions extracted from texts. To reach this goal it was necessary to construct datasets of evaluation systems, which were used to derive supervised machine learning models and compute

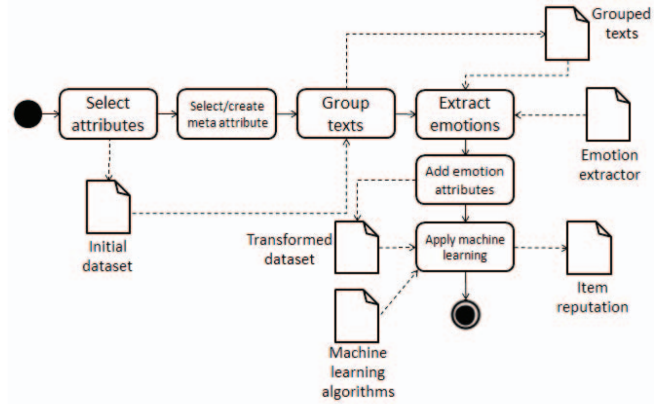

Fig. 1: Activity diagram of the model.

the correlations between emotions and numerical data (meta-attribute) presented on the datasets. In order to evaluate our approach, we used a text review corpus from Trip Advisor[2] originally extracted by [46] and other from Good Reads[3]. Good Reads dataset was extracted in September 2015 using the WebHarvy Web Scraper[4].

Trip Advisor is a website that presents user's opinions and evaluations about hotels, restaurants, and city attractions. Similarly, Good Reads is focused on user's opinions and evaluations about books. In both scenarios, users are allowed to provide textual reviews about items, in conjunction to other numeric data, e.g. in Trip Advisor, users can recommend or not a hotel setting a score, while in Good Reads users can also give scores to books, add it to their personal online shelf, or add it to a to-read list.

A text corpus is compatible with the proposed model if it presents two characteristics: (i) data about items and (ii) user's reviews about such items. It is important to emphasize that we generically consider as "items" something that is evaluated by users, such as hotels (Trip Advisor) or books (Good Reads). The main prerequisite is the existence of reviews, since they contain the texts from where the emotions will be extracted, as well as the numeric evaluation made by the users. In short, a review must present texts written about a given item and a corresponding numerical score representing its reputation.

Fig. 1 presents all the steps to compute reputation values. The first one is to **select attributes**, which aims at analyzing available data on the environment and gathers attributes that represent possibly useful information about the scenario. This process returns the initial dataset. To exemplify it, consider Tab. II, which presents the data extracted from Trip Advisor for each hotel, while Tab. III states initial data extracted from Good Reads.

The second step on Fig. 1 is to **select/create the meta attribute**, which is the construction of the numeric reputation model. This meta-attribute is used by supervised machine

[2]http://tripadvisor.com
[3]http://goodreads.com
[4]https://www.webharvy.com

TABLE I: Example of text review and emotion extraction.

Phrases	Result
"I came here with my boyfriend in January."	General weight: 0.00
"So convenient when you stay in NYC for a week and don't want to eat out every meal, every day."	General weight: 0.44 Valence: -1 (negative) Sadness weight: 0.19 Fear weight: 0.07
"Anyway, I'm a priority club member so they bumped me up to a room with a nicer view without asking :)."	General weight: 1.00 Valence: +1 (positive) Happiness weight: 1.00 Fear weight: 0.06

TABLE II: Trip Advisor initial data.

Field	Description
Number of reviews about the hotel	Total amount of reviews written about the hotel
Recommend yes	Total number of users who will recommend the hotel to friends
Recommend no	Total number of users who will not recommend the hotel to friends
Recommend undefined	Total number of users who did not answer if they would recommend or not
Hotel stars	Number of star the hotel has, ranging from 1 to 5 stars
Reputation	**(value rating + rooms rating + location rating + cleanliness rating + service rating + sleep rating + overall rating) / 7 - ranging from 0 to 1**

TABLE III: Good Reads initial data.

Field	Description
Score	Score the book automatically receives from the site
Votes	Total number of users who voted in the book
Total ratings	Total number of users who evaluated the book
Total reviews	Total number of reviews written about the book
Add person	Total number of people who added the book to their virtual shelves
To reads	Total number of people who wants to read the book
Reputation	**Rating users gave to the book, ranging from 0 to 1**

learning algorithms and reflects the reputation of items. Additionally, reputation can be manually created using the attributes extracted on the first step in the case that this information is not publicly available. Due to the fact that one of the dataset prerequisites is the existence of reviews, the meta-attribute will be often presented because reviews are composed by texts and corresponding numerical scores. Therefore, reputation (label of instances) is given by the average scores for each item, given by $\frac{1}{n}\sum_{s_i \in S} s_i$, where n is the amount of scores of an item and s_i iterates over all scores in the set of scores S of the same item.

To exemplify, in Trip Advisor's users rate the hotels in seven categories: value, rooms, location, cleanliness, service, sleep, and overall. Thus, the meta attribute is given by the average of these ratings, since each hotel can be evaluated by many users. Similarly, in Good Reads users set only one rate to the book, so the meta attribute is simply an average of this rate. Both attributes can be seen on Tabs. II and III, respectively.

The third step is to **group texts** for each item, and it receives as input texts from the initial dataset and returns grouped texts. This stage is necessary because we need to extract emotions from all texts written about an item, so all reviews must be clustered altogether.

Afterwards, the next step is to **extract emotions** from the grouped texts using an emotion extractor. Attributes returned by the emotion extractor must be added on a new dataset called **transformed dataset**, which contains the original data (Tabs. II and III) plus the emotions and valence. Tab. IV presents these new attributes.

The two positive emotions attributes are represented by the sum of positive emotions: happiness and surprise. On the other hand, the four attributes regarding negative emotions are the sum of negative emotions: sadness, fear, disgust, and anger. The remainder of attributes were explained in Sec. III.

The inclusion of these attributes relies on the emotion extractor, since other tools may return different data. It is important to emphasize that the proposed model is not dependent on specific tools, and other emotion extractors can be used instead of Synesketch [39].

Finally, the last step is to **apply machine learning** algorithms to construct the reputation model, receiving as input the transformed dataset and regression algorithms according to Fig. 1. This step will be detailed on Sec. VI.

Tab. V presents statistics about the data of both datasets. It is possible to observe that the hotel and book which contain the fewer number of reviews is only 1. We used the maximum quantity of data available on the datasets to avoid loss of potentially useful data. One can also see on Tab. V that for Trip Advisor dataset a number of about 15% of phrases was discarded by the emotion extractor because no emotions were found. This number increases to 25% in the Good Reads dataset.

VI. Analysis

In this section we access the performance of our proposal with several inductive machine learning algorithms in the task of predicting the reputation of an item given three datasets for each scenario. The first dataset D_1 presents only the numeric data and it is represented by Tables II and III, which corresponds to the Trip Advisor and Good Reads initial data, respectively. The second dataset, D_2, corresponds exclusively to the emotional attributes shown on Table IV. Finally, the third dataset D_3 is composed by the initial data plus the emotional attributes.

Based on our initial hypothesis, the expected result is that D_2 or D_3 present higher correlation coefficient when compared to D_1. Therefore, if D_2 or D_3 has higher correlation coefficient with statistically significant differences, it is possible to claim that emotions extracted from texts using Synesketch [39] impact on the guessing of items' reputation.

In the following sections we survey evaluated algorithms and state the experimental protocol adopted, thus enabling a discussion of the empirical results obtained.

A. Learning Algorithms

In order to test our hypothesis, the datasets described on the latter section are used to generate models given the following inductive learners. Algorithms were chosen due to their broad usage in the machine learning community, their different biases, and results obtained in different applications.

Linear Regression. Linear regression is a commonplace inductor due to its simplicity. It learns a linear function by the least squares fit and eliminates co-linear attributes present in the learning dataset [47].

M5P. M5P is a decision tree regression model [48] where linear models are located at the leaves and split nodes are created recursively by minimizing the intra-set variation in output values on each branch. Besides being powerful, M5P also generates models that are compact and relatively human comprehensible.

M5Rules. M5Rules generates and maintains a decision list for regression problems using a divide-and-conquer approach.

TABLE IV: Emotion attributes.

Emotion Attributes	
Happiness quantity	Positive emotions quantity
Happiness average weight	Positive emotions average weight
Sadness quantity	Negative emotions quantity
Sadness average weight	Negative emotions average weight
Fear quantity	Positive valence quantity
Fear average weight	Positive valence average
Anger quantity	Negative valence quantity
Anger average weight	Negative valence average
Disgust quantity	Neutral valence quantity
Disgust average weight	Neutral valence average
Surprise quantity	Phrase quantity
Surprise average weight	

Iteratively, M5Rules builds a model tree using M5P and makes the branch to the most homogeneous leaf a new rule [49].

SMOReg. SMOReg is a support vector machine for regression that works as a sequential minimal optimization algorithm using polynomial or radial basis function kernels and cost parameters [50].

B. Experimental Protocol

The utility of algorithms in experiments is given by the R^2 goodness-of-fit correlation coefficient. R^2 can be computed as the square of the correlation between the observed reputation y values and the predicted reputation $\hat{y}$ values. Alternatively, $R^2 \in [0;1]$ is given by Eq. 1, where summations are over all observations.

$$R^2 = \frac{\sum(\hat{y}_i - \overline{y})^2}{\sum(y_i - \overline{y})^2} \tag{1}$$

Predictions close to actual values imply in higher R^2 until its maximum bound of 1. Conversely, if predictions are completely unrelated to actual values, we have $R^2 = 0$. Our evaluation encompasses 50 runs of a 10-fold cross-validation, therefore, diminishing the probability of R^2 overfitting.

In order to verify if there is significant statistical differences between algorithms, we proceeded with a combination of Wilcoxon's signed ranks test or Friedman's and Bonferroni-Holm's non-parametric hypothesis tests [51], [52] with a confidence level of 95%. Finally, all algorithms implementations are given, and experiments were performed, under the Waikato Environment for Knowledge Analysis (WEKA) framework [47].

C. Discussion

In Tabs. VI and VII we present the results obtained in Trip Advisor and Good Reads experiments, respectively. In these results, one can see that D_3 presents higher R^2 values with all classifiers, in both Trip Advisor and Good Reads experiments. This is important since these results show that the addition of emotional features into our datasets enables higher predictive capabilities, therefore, boosting a reputation guessing agent's utility. One can observe that the correlation coefficient during the SMOReg in the Good Reads experiments were low, and we

TABLE V: Text corpus statistics.

Dataset	Trip Advisor	Good Reads
Number of reviews	$11,937$	$48,345$
Number of characters	$10,537,693$	$33,798,600$
Average characters per review	882	699
Number of phrases	$110,210$	$330,685$
Number of phrases with emotions	$93,250$	$248,797$
Number of items	619	932
Average of reviews per item	19	52
Average of phrases per item	150	266
Minimum number of reviews	1	1
Maximum number of reviews	104	177

TABLE VI: Correlations obtained during the Trip Advisor experiment.

R^2 – Trip Advisor			
Learner	D_1	D_2	D_3
Linear Regression	0.73	0.69	**0.77**
M5P	0.81	0.63	**0.83**
M5Rules	0.79	0.66	**0.82**
SMOReg	0.73	0.68	**0.78**

claim it happens because the predictive attributes are highly imbalanced for this dataset. Additionally, during the SMOReg experiments we tested only the polynomial kernel, thus, we argue that this kernel did not produce a good approximation between the values for this dataset, as one can see on Tab. VII. A future alternative to have better results is to perform more experiments using other kernels that might find superior linear separation.

In order to provide statistical foundation to our analysis, we proceeded with non-parametric statistical tests. With the aid of Friedman's and Bonferroni-Holm's tests, we were able to determine that all learners applied on D_3 outperformed D_1 and D_2 in both Trip Advisor and Good Reads datasets, both with $p < .05$. Again, these results confirm our initial hypothesis that through the embedding of emotions, an intelligent agent is capable of boosting its items reputation guessing accuracy significantly when compared to non-emotion-wise proposals.

On the other hand, results show that D_2 presents the average worst results. At a first glance, it may seem that emotions do not correlate at all with reputation guessing, yet, we show that differences are not as big in Trip Advisor experiment, while there is no significant difference between D_1 and D_2 in Good Reads experiments accordingly to Wilcoxon's signed rank test with $p < .05$. In practice, our empirical evidence highlights that the aggregation of both public available attributes (those provided from Trip Advisor and Good Reads) and emotional attributes leads to better overall results.

One could argue that the differences obtained between M5P and M5Rules with and without emotion features are small, and therefore, meaningless. In Fig. 2 we present partial trees obtained by M5P in the Good Reads experiment. By definition, higher levels of decision trees possess attributes that better characterize data. Therefore, one can see that in Fig. 2b, besides the "votes" attribute, the "positive emotions average weight" and "phrase quantity" attributes appear in higher levels of the tree, while the tree built without emotions differs significantly (Fig. 2a). Besides serving as split nodes in the higher levels of the trees, all the emotional attributes are also used in the linear regression models located at the leaves. In practice, this highlights that emotional attributes are significant to the task of guessing an item's reputation. We refrain from providing a partial tree for the Trip Advisor experiments since the same attributes appear in the three highest levels of the tree in all experiments. Another important trait that seems evident is the independence between the attributes. As M5P and M5Rules perform local search in the problem space, we claim that they can be more suitable when there is no strong correlation among the prediction variables.

TABLE VII: Correlations obtained during the Good Reads experiment.

R^2 – Good Reads			
Learner	D_1	D_2	D_3
Linear Regression	0.22	0.34	**0.42**
M5P	0.62	0.42	**0.65**
M5Rules	0.59	0.34	**0.62**
SMOReg	0.12	0.33	**0.35**

Finally, one can observe on Tabs. VI and VII that the best results were achieved with the M5P algorithm. However, although the results show significance statistical difference, they do not show much advantages toward the proposed approach against of just employing numerical data. We aimed to show with our approach that the use of emotions can validate our hypothesis that emotions are important to guessing items reputation, and we claim that the use of selection algorithm in D_3 will lead to much better results.

Finally, we emphasize that our proposal relies on emotion extractors, e.g. Synesketch [39], so if these tools return more robust results, our method is also expected to provide better results.

VII. Conclusion

In this paper we proposed an emotion bound reputation model based on Ekman's emotions which was tested in the scenario of evaluation systems. To validate and experiment our proposal we used data obtained from Trip Advisor and Good Reads evaluation systems websites. Empirical results obtained validate our initial hypothesis that emotion-based intelligent agents possess as good or better utility when compared to those that do not account for emotional characteristics with $p < .05$. In other words, these results fit the definitions of affective trust [6], since they showed improvements to guessing items' reputation. Our results show that there is the need of future research onto how to embed emotions in trust and reputation systems.

Future works include (i) an analysis on feature selection algorithms to improve agents' utility; (ii) the broadening of experiments, by encompassing others datasets from other evaluation systems, in order to corroborate the effectiveness of

(a) Partial tree obtained by D_1.

(b) Partial tree obtained by D_3.

Fig. 2: Partial trees obtained during Good Reads experiment.

our proposal in different scenarios; (iii) the usage of different emotion extractor APIs, as well as the ones based on the Appraisal Theory. The adoption of Appraisal Theory, particularly the OCC model [11], will lead to the same quantity of positive and negative emotions, differently from Ekman's models that has two positive and four negative emotions, which might generate a bias to the induction process. Other future works include (iv) the use of emotional features against and with different types of machine learning algorithms like bag-of-words, vectorial space and neural networks, (v) a comparison with state-of-the-art bag of words text mining techniques, and finally (vi) a comparison between other regression evaluation metrics aiming to evaluate the linearity of the relation, such as the maximal information coefficient [51].

Acknowledgment

This project was financially supported by the *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* (CAPES) through the *Programa de Suporte à Pós-Graduação de Instituições de Ensino Particulares* (PROSUP) program.

References

[1] J. Granatyr, V. Botelho, O. R. Lessing, E. E. Scalabrin, J.-P. Barthès, and F. Enembreck, "Trust and reputation models for multiagent systems," *ACM Comput. Surv.*, vol. 48, no. 2, pp. 27:1–27:42, Oct. 2015. [Online]. Available: http://doi.acm.org/10.1145/2816826

[2] A. S. Rao and M. P. Georgeff, "Bdi agents: From theory to practice," in *In Proceedings of the First International Conference on Multi-agent Systems (ICMAS-95*, 1995, pp. 312–319.

[3] R. Trappl, P. Petta, and S. Payr, *Emotions in Human and Artifacts*. Cambridge: MIT Press, 2003.

[4] P. Thagard, *Hot Tought: Mechanisms and Applications of Emotional Cognition*. Cambridge: MIT Press, 2008.

[5] M. A. S. N. Nunes, "Tendências à tomada de decisão computacional," in *Proceedings of the W3C*, São Paulo, Brazil, 2009.

[6] D. J. McAllister, "Affect and cognition-based trust as foundations for interpersonal cooperation in organizations," *The Academy of Management Journal*, vol. 38, no. 1, pp. 24–59, 1995.

[7] B. Neville and J. Pitt, "A computational framework for social agents in agent mediated e-commerce," in *Engineering Societies in the Agents World IV*, A. Omicini, P. Petta, and J. Pitt, Eds. Springer-Verlag, 2004. [Online]. Available: http://www.ai.univie.ac.at/~paolo/conf/ESAW03/esaw03accpapers.html

[8] A. Pinto, *Simulação e Avaliação de Comportamentos em Sistemas Multi-Agentes baseados em Modelos de Reputação e Interação*. São Paulo, Brazil: UNISINOS - Computação Aplicada (PIPCA) - RS/USP - ICMC, 2008.

[9] G. K. Bitencourt, R. A. Silveira, and J. Marchi, "Truste: An emotional trust model for agents," in *Proceedings the Workshop-Escola de Sistemas de Agentes, seus Ambientes e Aplicações (WESAAC)*, São Paulo, Brazil, 2013.

[10] V. Carofiglio and F. deRosis, "In favour of cognitive models of emotions," in *AISB'05 Workshop on Mind-Minding Agents*, 2005.

[11] A. Ortony, G. L. Clore, and A. Collins, *The cognitive structure of emotions*. Cambridge, New York: Cambridge Univ. Press, 1994, autre tirage: 1999. [Online]. Available: http://opac.inria.fr/record=b1125551

[12] P. Ekman, "An argument for basic emotions," *Cognition and Emotion*, pp. 169–200, 1992.

[13] R. Marimon, J. P. Nicolini, and P. Teles, "Competition and reputation," European University Institute, Economics Working Papers, 1999. [Online]. Available: http://EconPapers.repec.org/RePEc:eui:euiwps:eco99/18

[14] Plato, *The Republic (370BC)*. Viking Press, 1955.

[15] V. Buskens, "The social structure of trust," *Social Networks (20)*, pp. 265–298, 1998.

[16] D. Bromley, "Reputation, image, and impression management," *Psychology and Marketing*, vol. 12, no. 2, pp. 161–163, 1995. [Online]. Available: http://dx.doi.org/10.1002/mar.4220120207

[17] M. Karlins and H. I. Abelson, *Persuasion: how opinions and attitudes are changed*. Springer Pub. Co., 1959.

[18] D. Artz and Y. Gil, "A survey of trust in computer science and the semantic web," *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*, 2007.

[19] S. Castaldo, "Trust variety - conceptual nature, dimensions and typologies," in *IMP Conference Lugano*, 2003, pp. 4–6.

[20] J. A. Simpson, "Psychological foundations of trust," *Current Directions in Psychological Science*, pp. 264–268, 2007.

[21] D. H. McKnight, V. Choudhury, and C. Kacmar, "The impact of initial consumer trust on intentions to transact with a web site: a trust building model," *Journal of Strategic Information Systems*, vol. 11, no. 3-4, pp. 297–323, 2002.

[22] J. K. Butler, "Toward understanding and measuring conditions of trust: Evolution of a conditions of trust inventory," *Journal of Management 17(3)*, pp. 643–663, 1991.

[23] D. S. Johson and K. Grayson, *Sources and Dimensions of Trust in Service Relationships*. SAGE Publications, Inc., 2000, pp. 357–371. [Online]. Available: http://dx.doi.org/10.4135/9781452231327

[24] J. D. Lewis and A. J. Weigert, "Social atomism, holism, and trust," *Sociological Quarterly*, vol. 26, no. 4, pp. 455–471, 1985. [Online]. Available: http://dx.doi.org/10.1111/j.1533-8525.1985.tb00238.x

[25] C. Johnson-George and W. Swap, "Measurement of specific interpersonal trust: Construction and validation of a scale to assess trust in a specific other," *Journal of Personality and Social Psychology*, vol. 43, pp. 1306–1317, 1982.

[26] K. H. Hanzaee and A. Norouzi, "The role of cognitive and affective trust in service marketing: Antecedents and consequence," *Research Journal of Applied Sciences, Engineering and Technology*, vol. 4, pp. 4996–5002, 2012.

[27] Y. Wang and J. Vassileva, "Trust and reputation model in peer-to-peer networks." in *Peer-to-Peer Computing*, N. Shahmehri, R. L. Graham, and G. Caronni, Eds. IEEE Computer Society, 2003, pp. 150–157. [Online]. Available: http://dblp.uni-trier.de/db/conf/p2p/p2p2003.html#WangV03

[28] B. Misztal, *Trust in Modern Societies: The Search for the Bases of Social Order*. Polity, 1996.

[29] Z. Qing-hua, W. Chong-jun, and X. Jun-yuan, "Core: A trust model for agent coalition formation," in *Proceedings of the 2009 Fifth International Conference on Natural Computation - Volume 05*, ser. ICNC '09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 541–545. [Online]. Available: http://dx.doi.org/10.1109/ICNC.2009.612

[30] J. Bentahar and J. C. Meyer, "A new quantitative trust model for negotiating agents using argumentation," *IJCSA*, vol. 4, no. 2, pp. 1–21, 2007. [Online]. Available: http://www.tmrfindia.org/ijcsa/V4I21.pdf

[31] R. Hermoso, H. Billhardt, and S. Ossowski, "Role evolution in open multi-agent systems as an information source for trust," in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*, ser. AAMAS '10. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 217–224. [Online]. Available: http://dl.acm.org/citation.cfm?id=1838206.1838237

[32] S. Liu, H. Yu, C. Miao, and A. C. Kot, "A fuzzy logic based reputation model against unfair ratings," in *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, ser. AAMAS '13. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2013, pp. 821–828. [Online]. Available: http://dl.acm.org/citation.cfm?id=2484920.2485051

[33] A. Damasio, *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Quill, 1994.

[34] R. Lazarus, *Emotion and Adaptation*. Oxford University Press, 1991.

[35] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up?: Sentiment classification using machine learning techniques," in *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing - Volume 10*, ser. EMNLP '02. Stroudsburg, PA, USA: Association for Computational Linguistics, 2002, pp. 79–86. [Online]. Available: http://dx.doi.org/10.3115/1118693.1118704

[36] J. O'Donovan, B. Smyth, V. Evrim, and D. McLeod, "Extracting and visualizing trust relationships from online auction feedback comments," in *IJCAI*, 2007, pp. 2826–2831.

[37] A. Balahur, J. M. Hermida, and A. Montoyo, "Detecting implicit expressions of emotion in text: A comparative analysis," *Decis. Support Syst.*, vol. 53, no. 4, pp. 742–753, Nov. 2012. [Online]. Available: http://dx.doi.org/10.1016/j.dss.2012.05.024

[38] B. Martinazzo, M. M. Dosciatti, and E. C. Paraiso, "Identifying emotions in short texts for brazilian portuguese," in *IV International Workshop on Web and Text Intelligence*, Curitiba, Brazil, 2012, pp. 1–10.

[39] U. Krcadinac, P. Pasquier, J. Jovanovic, and V. Devedzic, "Synesketch: An open source library for sentence-based emotion recognition." *T. Affective Computing*, vol. 4, no. 3, pp. 312–325, 2013.

[40] S. M. Mohammad, "Sentiment analysis: Detecting valence, emotions, and other affectual states from text," *Emotion Measurement (forthcoming)*, 2015.

[41] B. Selvam and S. Abirami, "A survey on opinion mining framework," *International Journal of Advanced Research in Computer and Communication Engineering*, vol. 2, 2013.

[42] G. A. Miller, "Wordnet: A lexical database for english," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, Nov. 1995. [Online]. Available: http://doi.acm.org/10.1145/219717.219748

[43] M. Johns and B. G. Silverman, "How emotions and personality effect the utility of alternative decisions: A terrorist target selection case study," in *Proceedings of the Tenth Conference on Computer Generated Forces and Behavioral Representation, SISO*, 2001, pp. 55–64.

[44] J. R. Dunn and M. E. Schweitzer, "Feeling and believing: The influence of emotion on trust," *Journal of Personality and Social Psychology*, pp. 736–748, 2005.

[45] A. Sloman, "Motive mechanisms and emotions," in *Emotional and Cognition*, 1987, pp. 217–234.

[46] A. Roshchina, J. Cardiff, and P. Rosso, "User profile construction in the twin personality-based recommender system," in *IJCNLP Workshop on Sentiment Analysis where AI meets Psychology, SAAIP-2011, 5th Int. Joint Conf. on Natural Language Processing*, 2011.

[47] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The weka data mining software: An update," *SIGKDD Explor. Newsl.*, vol. 11, no. 1, pp. 10–18, Nov. 2009. [Online]. Available: http://doi.acm.org/10.1145/1656274.1656278

[48] Y. Wang and I. H. Witten, "Inducing model trees for continuous classes," in *In Proc. of the 9th European Conf. on Machine Learning Poster Papers*, 1997, pp. 128–137.

[49] G. Holmes, M. Hall, and E. Prank, "Generating rule sets from model trees," in *Advanced Topics in Artificial Intelligence*, ser. Lecture Notes in Computer Science, N. Foo, Ed. Springer Berlin Heidelberg, 1999, vol. 1747, pp. 1–12. [Online]. Available: http://dx.doi.org/10.1007/3-540-46695-9_1

[50] S. Shevade, S. Keerthi, C. Bhattacharyya, and K. Murthy, "Improvements to the smo algorithm for svm regression," *Neural Networks, IEEE Transactions on*, vol. 11, no. 5, pp. 1188–1193, Sep 2000.

[51] G. W. Corder and D. I. Foreman, *Nonparametric Statistics for Non-Statisticians: A Step-by-Step Approach*. New Jersey: Wiley, 2009.

[52] M. Friedman, "The Use of Ranks to Avoid the Assumption of Normality Implicit in the Analysis of Variance," *Journal of the American Statistical Association*, vol. 32, no. 200, pp. 675–701, Dec. 1937. [Online]. Available: http://dx.doi.org/10.2307/2279372