

## DETECTION OF FAKE ONLINE REVIEWS USING SUPERVISED AND SEMI SUPERVISED LEARNING

<sup>1</sup>MR.K.VAMSI KRISHNA, <sup>2</sup>CH.MINI, <sup>3</sup>B.DHANALAKSHMI, <sup>4</sup>U.SOUMYA

<sup>1</sup>Associate professor, Department of CSE, NRI Institute Technology, pothavarapadu(v) via NunnaAgiripalli(M), pin-521212.

<sup>2,3,4</sup> Student , NRI Institute Technology, pothavarapadu(v) via NunnaAgiripalli(M), pin-521212.

**Abstract**—In present era, in brand selection we mainly depend on online reviews and review sites are more confronted with the spread of misinformation, i.e., opinion spam, which aims at promoting or damaging some target businesses, by misleading either human readers, or automated opinion mining and sentiment analysis systems. For this reason, in the last years, several data-driven approaches have been proposed to assess the credibility of user-generated content diffused through social media in the form of on-line reviews. Distinct approaches often consider different subsets of characteristics, i.e., features, connected to both reviews and reviewers, as well as to the network structure linking distinct entities on the review-site in exam. This article aims at providing an analysis of the main review and reviewer-centric features that have been proposed up to now in the literature to detect fake reviews, in particular from those approaches that employ supervised machine learning techniques. These solutions provide in general better results with respect to purely unsupervised approaches, which are often based on graph-based methods that consider relational ties in review sites. Furthermore, this work proposes and evaluates some additional new features that can be suitable to classify genuine and fake reviews. For this purpose, a supervised classifier based on Random Forests have been implemented, by considering both well-known and new features, and a large-scale labeled dataset from which all these features have been extracted. The good results obtained show the effectiveness of new features to detect in particular singleton fake reviews, and in general the utility of this study.

### I. INTRODUCTION

The social Web and the increasing popularity of social media have led to the spread of multiple kinds of content (i.e., textual, acoustic, visual) generated directly by users, the so called *user-generated content* (UGC). By means of Web 2.0 technologies, it is possible for every individual to diffuse contents on social media, almost without any form of trusted external control. This implies that there are no means to verify, a priori, the reliability of the sources and the believability of the content generated. In this context, the issue of assessing the credibility of the information diffused by means of social media platforms is receiving increasing attention from researchers.

In particular, this issue has been deeply investigated in review sites, where the spread of misinformation in the form of *opinion spam*, and

the negative consequences that it brings, are particularly harmful for both businesses and users. In this context, opinion spam detection aims at identifying fake reviews, fake comments, fake blogs, fake social network postings, deceptions, and deceptive messages [1], and to make them readily recognizable. Detection techniques to identify *fake reviews* have been proposed in particular for specific review sites such as TripAdvisor<sup>1</sup> or Yelp, <sup>2</sup> where users' reviews have a powerful effect on people visiting the Website for advice. Therefore, a recommendation of a product or a service such as a restaurant or a hotel based on false information can have detrimental consequences.

Most approaches that have been proposed so far to detect fake reviews in these social media platforms rely on supervised machine learning

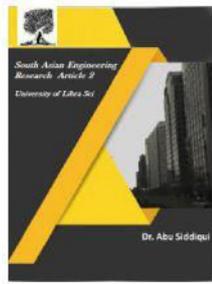


2581-4575

# International Journal For Recent Developments in Science & Technology



A Peer Reviewed Research Journal



techniques and on distinct characteristics, i.e., *features*, connected to the reviews and/or the reviewers who generated them. It has been shown in the literature that their usage can lead to an effective identification of suspicious contents and/or reviewers' behaviors, and consequently of misinformation [2].

Recent approaches have suggested the additional use of features that consider the social structure of the network underlying the considered review site. These approaches, which are often based on unsupervised graph based methods, usually provide worse performance with respect to supervised solutions. On the other hand, supervised approaches too present some issues. First, available solutions have often considered a small set of features, or distinct classes of features separately; second, they have been evaluated on small datasets extracted from the well-known review sites previously cited. Thus, the proposed solutions are in most of the cases partial, or review-site-dependent.

Considering the variety of features that have been proposed and used separately by supervised approaches, the goal of this article is to provide a *feature analysis* illustrating the most suited and general *review-* and *reviewer-centric features* that can be employed in the review site context to detect fake reviews. Among these features, some are well-known and taken from the literature, others are *new* and constitute a further contribution of the paper. To evaluate the utility of this set of features in classifying genuine and fake reviews, a *supervised classifier* based on a well-known machine learning technique has been implemented. With respect to the literature, a publicly available *large-scale* and *general dataset* from the Yelp.com review site has been considered. This allows to provide more significant results with respect to the contribution of each feature taken singularly and of group of features. In particular, the important contribution of a specific group of features is

analyzing

the credibility of the so called *singleton reviews* has emerged. The promising results obtained show the effectiveness and the possible utility of the feature analysis illustrated in this article.

## II. RELATED WORK

In the last few years, depending on the context, researchers have proposed many different approaches to tackle the issue of the assessment of the *credibility* of the information diffused through social media [2]. Historically, the concept of credibility has been in turn associated with believability, trustworthiness, perceived reliability, expertise, accuracy, and with numerous other concepts or combinations of them [3].

According to Fogg and Tseng [4], credibility is a *perceived quality* of the *information receiver*, and it is composed of multiple dimensions. Different characteristics can be connected to: (i) the *source* of information, (ii) the *information* itself, i.e., its structure and its content, and (iii) the *media* used to disseminate information [5]. It has been demonstrated that, when considering these characteristics in terms of credibility, the impact of the delivery medium can change the perception that people have about sources of information and information itself [3], [5]. For this reason, one important question to be tackled nowadays is whether new media in the digital realm introduce new factors that may concur to credibility assessment [6],[7].

In the Social Web, evaluating information credibility deals with the analysis of the user-generated content [8], the authors' characteristics, and the intrinsic nature of social media platforms, i.e., the social relationships connecting the involved entities. These characteristics, namely *features*, can be simple *linguistic features* associated with the text of the UGC, they can be additional *meta-data features* associated for example with the content of a review or a tweet, they can also be extracted

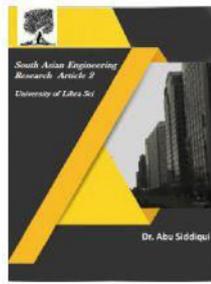


2581-4575

# International Journal For Recent Developments in Science & Technology



A Peer Reviewed Research Journal



from the behavior of the users in social media, i.e., *behavioral features*, or they can be connected to the user profile (if available). Furthermore, different approaches have taken into consideration *product-based features*, in the case of review sites where products and/or services are reviewed, or have considered *social features*, which exploit the network structure and the relationships connecting entities in social media platforms [9],[10].

In the last years, several approaches have been proposed to assess in an automatic or semi-automatic way the credibility of information in

this paper, the approaches that have produced the best results are generally based on supervised or semi-supervised machine learning techniques that take into account both review- and reviewer centric features. The first approaches were purely *linguistic*, in the sense that they employed simple textual features extracted from the text of reviews, often in the form of unigrams and/or bigrams [13], [14], [15], [16]. Other linguistic approaches have proposed generative classifiers

For this reason, more effective *multi-feature-based* approaches have been proposed, which employ several features of different nature in addition to simple linguistic ones, either by applying supervised or semi-supervised machine learning [1], [19], [21], or by implementing the Multi-Criteria Decision Making (MCDM) paradigm [22]. These approaches usually focus on small labeled datasets for evaluation purposes, constituted in most of the cases by ‘near ground truth’ data [9].

They usually avoid to consider features that are extracted from the social ties constituting the network of entities (e.g., users, products, reviews) considered by the review site. On the contrary, this kind of features is often utilized (together with the other features previously described) by graph-based approaches [23], [24]. These latter approaches are in most of the cases unsupervised, even if sometimes they can be

the Social Web; in particular, the most investigated tasks have been the identification of: (i) *opinion spam* in review sites [9], (ii) *fake news* in microblogging sites [11], and (iii) potentially harmful/inaccurate *online health information* [12]. In general, the majority of these approaches focus on data-driven techniques, which classify UGC with respect to credibility by employing different models.

With regard to opinion spam detection, and in particular to fake review detection, which is the focus of

based on language models [17], [18]. It has been demonstrated by Mukherjee et al. in [19] that focusing only on linguistic features is not effective to detect fake reviews from real datasets, since it is practically impossible for a human reader to distinguish between credible and not credible information by simply reading it, especially in an era where the skills in writing false reviews are constantly improving [20].

coupled with a supervised learning phase on a limited number of classification labels [25]. With respect to supervised approaches, totally unsupervised solutions generally provide slightly worst results [2], [9], [20].

This paper, by considering the effectiveness of supervised solutions, discusses and analyzes on a general level the most appropriate review- and reviewer-centric features that have been proposed so far in the literature to detect fake reviews; moreover, it proposes some new features suitable for this aim, in particular to detect singleton fake reviews, an issue that has not yet received the attention it deserves. To avoid the problem of the limited size of the labeled datasets considered up to now by the literature, two large-scale publicly-available datasets presented in [25] have been employed for evaluation purposes.

### III. FEATURE ANALYSIS

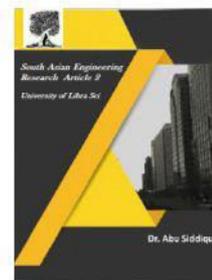


2581-4575

# International Journal For Recent Developments in Science & Technology



A Peer Reviewed Research Journal



As briefly introduced in Section II, many and different are the features that have been considered so far in the review site context to identify fake reviews. In some cases, features belonging to different classes have been considered separately by distinct approaches. In other cases, the employed features constitute a subset of the entire set of features that could be taken into account; furthermore, new additional features can be proposed and analyzed to tackle open issues not yet considered, for example the detection of singleton fake reviews. For these reasons, in this section we provide a global overview of the various features that can be employed to detect fake reviews. Both significant features taken from the literature and new features proposed in this article are considered. Since the most effective approaches discussed in the literature are in general supervised and consider *review-* and *reviewer-centric features*, these two classes will be presented in the following sections. The choices behind the selection of the features belonging to the above mentioned classes will be detailed along each section. When the features are taken from the literature, they will be directly referred to the original paper where they have been initially proposed. The absence of the reference will denote those features that have been widely used by almost every proposed technique.

Finally, the presence of the label denoted by [new] will indicate a feature proposed for the first time in this article. *A. Review-centric Features* The first class of features that have been considered, is constituted by those related to a *review*. They can be extracted both from the text constituting the review, i.e., *textual features*, and from *meta-data* connected to a review, i.e., *meta-data features*. In every review site, the *time* information regarding the publication of the review, and the *rating* (within some numerical interval) about the reviewed business are metadata, are always provided. In addition, in

relation to metadata features, those connected to the *cardinality* of the reviews written by a given user must be carefully studied. In fact, a large part of reviews are *singletons*, i.e., there is only one review written by a given reviewer in a certain period of time (this means that in the useraccount there is only one review at the time of the analysis). For this kind of reviews, specific features must be designed. In fact, as it will be illustrated in the following, many of the features that have been proposed in the literature are based on some statistics over several reviews written by the same reviewer. In the case of singletons, these features lose their relevance in assessing credibility.

Therefore, the definition of suitable features that are effective for detecting also singleton fake reviews becomes crucial. *1) Textual Features:* as briefly illustrated in Section II, it is practically impossible to distinguish between fake and genuine reviews by only reading their content. The analysis provided by Mukherjee et al. in [19] has shown that the KL-divergence between the languages employed by spammers and non spammers in Yelp is very subtle. However, the good results obtained in [26] by using linguistic features on a domain specific dataset (i.e., a Yelp's dataset containing only New Yorkjapanese restaurants), show that at least on a domain specific level, textual features can be useful. It is possible to use Natural Language Processing techniques to extract *simple* features from the text, and to use as features some *statistics* and some *sentiment estimations* connected to the use of the words.

- *Text:* several approaches employ as textual features both *unigrams* and *bigrams* extracted from the text of reviews, as illustrated in Section II.

- *Text statistics:* several statistics on the review content have been proposed as features by Li et al. in [21]:

- *Number of words*, i.e., the length of the

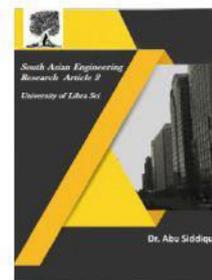


2581-4575

# International Journal For Recent Developments in Science & Technology



A Peer Reviewed Research Journal



review in terms of words;

- *Ratio of capital letters*, i.e., the number of words containing capital letters with respect to the total number of words in the review;
- *Ratio of capital words*, i.e., considering the words where all the letters are uppercase;
- *Ratio of first person pronouns*, e.g., 'I', 'mine', 'my', etc.;
- *Ratio of 'exclamation' sentences*, i.e., ending with the symbol '!'.

• *Sentiment evaluations:*

- *Subjectivity*, i.e., a number representing the proportion of subjective words (expressing sentiment, judgment) as opposed to objective (descriptive) words.

2) *Meta-data Features:* these kinds of features are extracted from the meta-data connected to reviews, or they can be generated by reasoning on the reviews' cardinality with respect to the reviewer and the entity reviewed.

• *Basic features:*

- *Rating*, i.e., the rating attributed in the review to the entity, in the form of some numerical value belonging to a given interval (e.g., 1-5 'stars');
- *Rating deviation* [27], i.e., the deviation of the evaluation provided in the review with respect to the entity's average rating;
- *Singleton* [25], i.e., it indicates the fact that the review is the only one provided by a reviewer in a given period of time (e.g., a day).

These basic features rely on some simple and intuitive heuristics. A fake review tends to contain a more 'extreme' rating with respect to genuine reviews, thus implying that the rating deviation from the entity's average rating is higher; furthermore, a singleton review provided by a user could indicate that s/he is not particularly involved in the review site community, which constitutes a possible indication of unreliability.

• *Burst features:* it is said that reviews for an entity are 'bursty' when there is a sudden concentration of reviews in a time period. These *review bursts* can be either due to sudden popularity of the entities reviewed or to spam attacks. Since it has been proven that reviews in the same

burst tend to have the same nature [28], it is possible to easily identify groups of fake reviews by analyzing the nature of the burst. Two burst detection studies have been described in [27], [28]. Taking inspiration from the just cited works, in this paper several features considering *burstiness* have been introduced. These features are related to the time window in which a review has been posted, relatively to a given reviewed entity. Basically, a review is more likely to be fake if it is posted on a day when the number of reviews is abnormally high, and when the average rating associated with an entity in a review (in a specific time window) varies significantly with respect to the entity's average rating (in general it decreases, for example passing from 3.5/5 to 2/5).

This assumption is valid for every kind of review, but it has proved to be particularly effective in boosting the detection of singleton fake reviews, which is particularly difficult without considering burstiness. With respect to the literature, the following new features have been considered:

- *Density [new]*, i.e., the number of reviews for a given entity on a given day of publication;
- *Mean rating deviation [new]*, i.e., the deviation of the average rating of an entity on the considered day with respect to the entity's average rating (in general);
- *Deviation from the local mean [new]*, i.e., the numerical evaluation if the rating assigned to a given entity in a review is close to the average rating assigned on the considered day;
- *Early time frame* [27], i.e., the amount of time it takes before the first review on a given

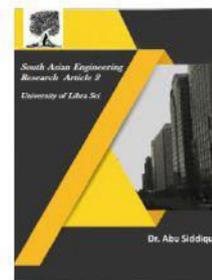


2581-4575

# International Journal For Recent Developments in Science & Technology



A Peer Reviewed Research Journal



entity is posted. In fact, often spammers review early to increase the impact of their (false) opinions on the audience.

## B. Reviewer-centric Features

This group of features is composed of features related to the reviewers' *behavior*. This way, it is possible to go beyond the content and meta-data associated with a review, which are limited for classification, and considering the behavior of users in general in writing reviews.

1) **Textual Features:** when considering reviewers' behavior, textual features are employed to address the problem of review duplication, which has been mentioned and studied in several research papers [1], [29], [30]. Specifically, the following textual features have been taken from [31]:

- *Maximum Content Similarity (MCS)*, i.e., the evaluation of the maximum similarity over the user's reviews;
- *Average Content Similarity (ACS)*, i.e., the evaluation of the average similarity over the user's reviews;
- *Word number average*, i.e., the average number of words that the user utilizes in her/his reviews. Users who duplicate their reviews will have a really high MCS. This can be considered as a suspicious behavior. Content similarity is expressed by means of the *cosine similarity* between the bag-of-words representations of the user's reviews.

2) **Rating Features:** they are based on some aggregation, for each considered reviewer, of the information concerning the ratings and/or the reviews s/he provided:

- *Total number of reviews*;
- *Ratios* [27], i.e., the ratio of negative, positive, and 'extreme' reviews (i.e., whose rating corresponds to the extremes of the considered rating interval);
- *Average deviation from entity's average* [31], i.e., the evaluation if a user's ratings assigned in her/his reviews are often very

different from the mean of an entity's rating (far lower for instance);

- *Rating entropy* [25], i.e., the entropy of rating distribution of user's (entity's) reviews;
- *Rating variance [new]*, i.e., the squared deviation of the rating assigned by a user with respect to the ratings mean. The variance as a rating feature has been added to further describe how the ratings for a particular user are distributed.

3) **Temporal Features:** they are based on the temporal information that further describes how the ratings are distributed over the time:

- *Activity time of the user* [31], i.e., the difference of timestamps of the last and first reviews for a given reviewer;
- *Maximum rating per day* [25], i.e., the maximum rating provided by a reviewer in the considered day;
- *Date entropy* [25], i.e., the temporal gap in days between consecutive pairs or reviews;
- *Date variance [new]*, i.e., the squared deviation of the timestamps in which a user post her/his reviews with respect to the timestamps mean. The variance as a temporal feature has been added to further describe how the reviews for a particular user are distributed over time windows.

## IV. SUPERVISED CLASSIFICATION

The classifier has been implemented to evaluate the impact of distinct features and set of features on the classification, and the overall performance that can be achieved. This constitutes an important part of the work, as it shows the impact of several kinds of features on a large-scale dataset, whereas many papers present only a small subset of these features, and evaluate them on small or review-site-dependent datasets, making it hard to assess their relative importance.

### Conclusion

In the last years, the issue of how to assess the

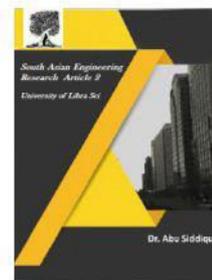


2581-4575

# International Journal For Recent Developments in Science & Technology



A Peer Reviewed Research Journal



credibility of social media information and in particular how to identify opinion spam in review sites has received increasing attention by researchers. Usually, the approaches to fake review detection are based on data-driven methods that consider several features associated with reviews, reviewers, and the network structure of the social network that can be used to classify reviews in terms of their credibility.

Supervised classifiers are in general more effective, and usually employ reviewer- and reviewer-centric features. Unsupervised classifiers are in most of cases based on graph-based models, and focuses on the social ties underlying the review site in exam (together with other kinds of features). Unsupervised solutions are in general less effective, but have the advantage that they do not need labeled datasets for training. Supervised solutions, on the contrary, have proven their effectiveness with respect to too small or review-site-dependent labeled datasets, and with respect to small subsets of features among the ones that have been proposed in general in the literature.

In this article, focusing on the effectiveness of supervised classification, a feature analysis has been performed, in order to summarize the main reviewer- and reviewer-centric features that are suited for fake review detection, and to propose new features that can be particularly useful to detect singleton reviews. To evaluate the impact of these features, a supervised classifier based on Random Forests has been developed.

To avoid the issues connected to the limited volume of available ground truths, a publicly available large-scale and general labeled dataset has been employed for evaluation purposes. The promising results obtained witness the utility of the proposed study.

## REFERENCES

[1] N. Jindal and B. Liu, "Opinion spam and analysis," in *Proceedings of the 2008 International Conference on Web Search and*

*Data Mining*". ACM, 2008, pp.219–230.

[2] M. Viviani and G. Pasi, "Credibility in Social Media: Opinions, News, and Health Information - A Survey," *WIREs Data Mining and Knowledge Discovery*, 2017. [Online]. Available: <http://dx.doi.org/10.1002/widm.1209>

[3] C. S. Self, "Credibility", in *An Integrated Approach to Communication Theory and Research, 2nd Edition*, M. B. Salwen and D. W. Stacks, Eds. Routledge, Taylor and Francis Group, 2008, pp. 435–456. [Online]. Available: <http://dx.doi.org/10.4324/9780203887011>

[4] B. J. Fogg and H. Tseng, "The elements of computer credibility," in *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 1999, pp.80–87.

[5] M. J. Metzger, A. J. Flanagin, K. Eyal, D. R. Lemus, and R. M. McCann, "Credibility for the 21st century: Integrating perspectives on source, message, and media credibility in the contemporary media environment," *Annals of the International Communication Association*, vol. 27, no. 1, pp. 293–335, 2003.

[6] M. J. Metzger and A. J. Flanagin, "Credibility and trust of information in online environments: The use of cognitive heuristics," *Journal of Pragmatics*, vol. 59, Part B, no. 0, pp. 210 – 220, 2013, biases and constraints in communication: Argumentation, persuasion and manipulation.

[8] M.-F. Moens, J. Li, and T.-S. Chua, Eds., *Mining User Generated Content*, ser. Social Media and Social Computing. Chapman and Hall/CRC, 2014.

[9] A. Heydari, M. ali Tavakoli, N. Salim, and Z. Heydari, "Detection of review spam: A survey," *Expert Systems with Applications*, vol. 42, no. 7, pp. 3634–3642, 2015.

[10] B. Carminati, E. Ferrari, and M. Viviani, "A multi-dimensional and event-based model for trust computation in the social web," in *International Conference on Social Informatics*.

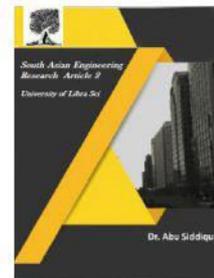


2581-4575

# International Journal For Recent Developments in Science & Technology



A Peer Reviewed Research Journal



- Springer, 2012, pp.323–336.
- C. Castillo, M. Mendoza, and B. Poblete, “Predicting information credibility in time-sensitive social media,” *Internet Research*, vol. 23, no. 5, pp. 560–588,2012.
- [8] T. J. Ma and D. Atkin, “User generated content and credibility evaluation of online health information: A meta analytic study,” *Telematics and Informatics*,2016.
- [9] K.-H. YooandU. Gretzel, “Comparison of deceptive and truthful travel reviews,” *Information and communication technologies in tourism 2009*, pp. 37–47,2009.
- [10] M. Ott, Y. Choi, C. Cardie, and J. T. Hancock, “Finding deceptive opinion spam by any stretch of the imagination,” in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies- Volume 1*. Association for Computational Linguistics, 2011, pp.309–319.
- [11] S. Banerjeeand A. Y. Chua, “Applauses in hotel reviews: Genuine or deceptive?” in *Science and Information Conference (SAI), 2014*. IEEE, 2014, pp.938–942.
- [12] D. H. Fusilier, M. Montes-y G´omez, P. Rosso, and R. G. Cabrera, “Detection of opinion spam with character n- grams,” in *International Conference on Intelligent Text Processing and Computational Linguistics*. Springer, 2015, pp.285–294.
- [13] C. L. Lai, K. Q. Xu, R. Y. K. Lau, Y. Li, and L. Jing, “Toward a language modeling approach for consumer review spam detection,” in *2010 IEEE 7th International Conference on E-Business Engineering*, Nov 2010,pp.1–8.
- [14] R. Y. K. Lau, S. Y. Liao, R. C.-W. Kwok, K. Xu, Y. Xia, and Y. Li, “Text mining and probabilistic language modeling for online review spam detection,” *ACM Trans. Manage. Inf. Syst.*, vol. 2, no. 4, pp. 25:1–25:30, Jan. 2012. [Online]. Available: <http://doi.acm.org/10.1145/2070710.2070716>
- [16] toappear.
- [15] A. Mukherjee, V. Venkataraman, B. Liu, and N. S. Glance, “What Yelp fake review filter might be doing?”in *Proceedings of ICWSM*,2013.
- [16] M. Crawford, T. M. Khoshgoftaar, J. D. Prusa, A. N. Richter, and H. Al Najada, “Survey of review spam detection using machine learning techniques,” *Journal of Big Data*, vol. 2, no. 1, p. 23,2015.
- [Online]. Available: <http://dx.doi.org/10.1186/s40537-015-0029-9>
- [17] F. Li, M. Huang, Y. Yang, and X. Zhu, “Learning to identify review spam,” in *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, vol. 22, no. 3, 2011, p.2488.
- [11] M. Viviani and G. Pasi, “Quantifier guided aggregation for the veracity assessment of online reviews,” *International Journal of Intelligent Systems*,2016.
- [12] G. Wang, S. Xie, B. Liu, and S. Y. Philip, “Review graph based online store review spammer detection,” in *2011 IEEE 11th International Conference on Data Mining*. IEEE, 2011, pp.1242–1247.
- [13] L. Akoglu, R. Chandy, and C. Faloutsos, “Opinion fraud detection in online reviews by network effects,” *ICWSM*, vol. 13, pp. 2–11,2013.
- [14] S. Rayana and L. Akoglu, “Collective opinion spam detection: Bridging review networks and metadata,” in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2015, pp. 985–994.
- [15] J. Fontanarava, G. Pasi, and M. Viviani, “An ensemble method for the credibility assessment of user-generated content,” in *WI’17 Proceedings - 2017 IEEE/WIC/ACM International Conference on Web Intelligence*, 2017,



2581-4575

# International Journal For Recent Developments in Science & Technology



A Peer Reviewed Research Journal

