

Mining micro-influencers from social media posts

Original

Mining micro-influencers from social media posts / Leonardi, Simone; Monti, DIEGO MICHELE; Rizzo, Giuseppe; Morisio, Maurizio. - ELETTRONICO. - (2020), pp. 867-874. (Intervento presentato al convegno SAC '20: The 35th ACM/SIGAPP Symposium on Applied Computing tenutosi a Brno (CZ) nel March 30-April 3, 2020) [10.1145/3341105.3373954].

Availability:

This version is available at: 11583/2773692 since: 2020-06-11T15:18:46Z

Publisher:

ACM

Published

DOI:10.1145/3341105.3373954

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

ACM postprint/Author's Accepted Manuscript

(Article begins on next page)

Mining Micro-Influencers from Social Media Posts

Simone Leonardi
Politecnico di Torino
Turin, Italy
simone.leonardi@polito.it

Giuseppe Rizzo
LINKS Foundation
Turin, Italy
giuseppe.rizzo@linksfoundation.com

Diego Monti
Politecnico di Torino
Turin, Italy
diego.monti@polito.it

Maurizio Morisio
Politecnico di Torino
Turin, Italy
maurizio.morisio@polito.it

ABSTRACT

Micro-influencers have triggered the interest of commercial brands, public administrations, and other stakeholders because of their demonstrated capability of sensitizing people within their close reach. However, due to their lower visibility in social media platforms, they are challenging to be identified. This work proposes an approach to automatically detect micro-influencers and to highlight their personality traits and community values by computationally analyzing their writings. We introduce two learning methods to retrieve Five Factor Model and Basic Human Values scores. These scores are then used as feature vectors of a Support Vector Machines classifier. We define a set of rules to create a micro-influencer gold standard dataset of more than two million tweets and we compare our approach with three baseline classifiers. The experimental results favor recall meaning that the approach is inclusive in the identification.

CCS CONCEPTS

• **Theory of computation** → **Machine learning theory**; • **Computing methodologies** → **Natural language processing**; **Lexical semantics**; • **Applied computing** → *Psychology*; *Economics*; *Sociology*;

KEYWORDS

Micro-influencer, Big5, Schwartz Values Theory, Natural Language Processing, Twitter

ACM Reference Format:

Simone Leonardi, Diego Monti, Giuseppe Rizzo, and Maurizio Morisio. 2020. Mining Micro-Influencers from Social Media Posts. In *The 35th ACM/SIGAPP Symposium on Applied Computing (SAC '20)*, March 30–April 3, 2020, Brno, Czech Republic. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3341105.3373954>

1 INTRODUCTION

Viral contents spread all around the globe with the help of social media platforms. Most of them are considered junk or useless, but they can also be exploited to convey positive messages. People

with influence power are able to encourage commonly accepted behaviours and contaminate large population slices.

In literature, people with influence power who are also active in social media platforms are defined influencers [9]. They tend to persuade other people to change their mind about topics and to modify their behaviors. Existing researches mainly analyse influencers in terms of notoriety and coverage of their posts to find them. These approaches try to detect candidates through social media platform scores by measuring followers reaction at certain posts (i.e. comments, sharing) and by inspecting social media graphs [1, 3, 11]. At the same time, they do not consider how the influence mechanism works at psychological level.

This work proposes an approach to automatically detect micro-influencers and to highlight their personality traits and community values by computationally analyzing their writings. Micro-influencers are a special kind of influencers, who are harder to find, less famous but with a higher engagement power over their communities [22]. The process starts from collecting user writings about trending topics. After a filtering and processing phase, we create the micro-influencers gold standard dataset following a rule-based approach.

The application of the lexical hypothesis [10], thanks to the vector space representation of words in the form of embeddings, allows the computation of personality traits and community-based values scores on micro-influencer writings. On a word level analysis, micro-influencer characteristics emerge from what the user writes in social media posts. We show the potential in approaching the influencing mechanism on a community based level through the work of Schwartz et al. [24] and on personality traits level through the work of McCrae and Costa [17] applied to the field of micro-influencing. Five Factor Model and Basic Human Values scores predicted with these methods are used as feature vectors of a Support Vector Machines classifier.

We have defined three research questions:

- RQ1: Which gold standard can be used for mining micro-influencers from social media posts?
- RQ2: How to extract their personality traits and community based values from text?
- RQ3: How to classify micro-influencers using a feature vector composed of personality traits and community based scores?

The remainder of this paper is structured as follows: in Section 2 we review related work, while in Section 3 we describe the proposed approach for detecting micro-influencers. In Section 4 we present a

SAC '20, March 30–April 3, 2020, Brno, Czech Republic

© 2020 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *The 35th ACM/SIGAPP Symposium on Applied Computing (SAC '20)*, March 30–April 3, 2020, Brno, Czech Republic, <https://doi.org/10.1145/3341105.3373954>.

preliminary experimental evaluation of our method. In Section 5 we discuss the obtained results and improvements. Finally, in Section 6, we conclude with a summary of the experiment outcome and outlining future work.

2 RELATED WORK

2.1 Micro-influencer metrics

Several studies used metrics to measure influencing abilities of individuals on social media. In [1], Anger et al. developed a score system called SNP (Social Network Potential) to identify top influencers in Austria. A similar research study was conducted by Bakshy et al. [3] who determined how influence analysis directly impact marketing investments. Furthermore, in [4], Bigonha et al. analysed social media posts to discover if users were talking about a certain topic and to understand if they were detractors or evangelists.

Finally, in [11], Kiss and Bichler created other scores to label influencer in social media. They described influence as a virus infecting other members inside a community having strict relationship with the disease origin. With a similar strategy, in [13], Leskovec et al. adopted centrality and diffusion as metrics related to the structure of relationships in social media friendship graph. This approach has common aspects to the one developed by Easley and Kleinberg in their analysis *Networks, Crowds and Markets* [7] where they wrote about cascading behaviours and disease spreading. In the work of Eliacik et al. [8], micro blogging were inspected. They developed metrics to find which users are central and act as influencers in the connected community, independently of the micro blogging topic.

Even if we are inspired by these previous works, we focus on metrics related to the niche case of micro-influencers. In Section 3 we explain how we measure and filter users with in-house rules.

2.2 NLP and Deep Learning

A different point of view is presented by studies that analyse influence from a linguistic perspective. In this field, Chen et al. [6] studied word use in social media to see correlation with community based values on Reddit. Kumar et al. [12] also worked on extracting personality traits and community values from written text. They dealt with the psycho-sociological compositions of social network community. They performed a massive analysis on corpus of text from Facebook, Twitter, collection of essay, correlated with personality assessment questionnaire and social network structures. Pennebaker et al. [19, 20] demonstrated the validity of personality traits extraction from text.

Concerning the usefulness of this approach on social media, Weisbuch et al. [26] showed how the average user tends to be spontaneous in writing posts without overthinking about what posting and thus allowing a correct personality detection.

Algorithms adopted in these research areas evolved through the years since 2003. Initially, just SVM with lexical features were adopted [2]. Lexical features such as the Appraisal lexical taxonomy is used by Argamon et al. [2] to detect the Neuroticism personality trait. Lexical taxonomy are hierarchical structures that follow linguistic rules to classify words. Homonymy, antonymy, synonymy, hyponymy and hypernymy define how words are related.

More recently, many research groups used deep learning due to the presence of increasing amount of available data and greater




Figure 1: Process pipeline. Starting from users and tweets retrieval querying Twitter, an initial filter on number of followers is applied. We then manipulate tweet objects both for text mining and for metadata analysis. Text is used in personality traits and community values extraction, while metadata in the gold standard creation. Finally, personality traits and community values feed a supervised classifier that has the micro-influencer labels, from gold standard, as expected output.

computational resources as in the study of Su et al. [25]. In this case, LIWC (Linguistic Inquiry Word Count) allows grammar annotation and Su applied it on text extracted from dialogues. The adoption of deep learning methods to process natural language avoids the definition of a previous set of lexical rules. Deep learning transforms words in a vector space thanks to the analysis of words distribution inside documents. Once transformed, words gain mathematical properties, such as similarity, that are exploited to understand how two words are similar computing the angle cosine between them.

Later, Majumder et al. [14] used a CNN to derive a fixed-length feature vector starting from word2vec word embeddings, which they extended with eighty four additional features from Mairesse’s library [18]. For classification, the so computed document vectors were fed both to a multi-layer perceptron (MLP) and to a polynomial SVM classifier.

In our work, we perform a regression with SVM for personality traits and we develop a different algorithm for community based values. It acts as a gravitational field in clustering words that belong to the same semantic area. We experiment many different classifiers to map micro-influencers and their associated traits and values.

On the use of word embeddings to classify sentences, another study by Mandelbaum and Shalev [15] obtained results suggesting that pre-trained vectors are universal feature extractors and they can be utilized across datasets. Our work uses GloVe [21]. GloVe provides pre-trained word embeddings created starting from a Wikipedia corpus. In Section 4 we describe how we apply it to personality traits and community-based values. We use embeddings as a vocabulary to establish the spatial word representation. The importance and diversification of various approach in word embeddings was described by Roy Schwartz et al. in [23], where they proposed an improved similarity prediction method and also defined the distinction between words similarity and association. They highlighted how hyponymy, meronymy and antonym contribute to word clustering.

3 DETECTING MICRO-INFLUENCERS

In the following, we formalize the set of rules exploited to create the micro-influencer gold standard definition. We then explain how we use the lexical hypothesis to define an algorithm that extracts personality traits and community based values from user writings in the specific case of micro-influencer field.

3.1 Dataset creation and gold standard

We establish rules to create a gold standard in the form of an annotated dataset in the field of micro-influencers. We deal with the entire pipeline: data retrieval, label definition through our oracle and final dataset creation.¹

Rules defined on collected data are described below. We applied these rules to candidates having a number of followers in a range between 1k and 100k. These scores will explain how our oracle labels a user as a micro-influencer in the gold standard. We compute the average Embeddedness among all potential micro-influencers. We do the same for Interaction and Engagement. These averages act as threshold. If a potential micro-influencer has each of the three scores over the relative threshold, he is labeled as micro-influencer. Embeddedness, Interaction and Engagement are equally important.

Embeddedness score is derived from Easley et al. [7], where the authors talk about *neighborhood overlap*. Information spreads quickly inside a community where at least two influencing members speak about similar concepts. If two potential micro-influencers have almost the same followers, they belong to the same community. When we look at the list of followers of two micro-influencers talking about the same topic (i.e. *Artificial Intelligence*) we notice many followers of the first micro-influencer also appear as followers of the second one. This is equivalent to have more than two micro-influencers writing about the same topic.

For this reason, we reward a micro-influencer if his followers appear in many list of followers of the other micro-influencers. A larger split of audience means larger influence capability in proportion to the number of total micro-influencers in that area. At the same time, a single message written by this micro-influencer will be reinforced by many other micro-influencers, talking about the same topic, and sharing his followers.

¹The dataset containing source tweets, related computed scores and our gold standard is available at <https://doi.org/10.6084/m9.figshare.11309669>.

$$Embeddedness_i = \frac{\sum_j |\{Follower_i\} \cap \{Follower_j\}|}{NumberofFollowers_i} \quad (1)$$

In the previous equation, i is the potential micro-influencer we are analysing and j represents others potential micro-influencers in our list.

Interaction score is computed in Equation 2. We select for each tweet t how many retweets are performed by followers. We take inspiration from the *Interactor Ratio* of Anger and Kittl [1]. It is important to notice that people who can retweet but do not follow the potential micro-influencer are not considered.

A micro-influencer must increase the loyalty of his followers. He must interact with them and in return his followers will share his opinions. We are focused on his followers. A micro-influencer also persuades occasional users, but we do not consider them, because we want to know if his usual circle of followers interacts with him.

$$Interaction_i = \frac{\sum_t RetweetByFollower_{i,t}}{TotalFollowers_i \times NumberofTweets_i} \quad (2)$$

In Equation 2 we normalize the score on total number of tweets, otherwise a micro-influencer that has already written lot of tweets has an advantage with respect to last arrived micro-influencers.

Engagement score is modified from Grin tool.² Engagement is measured summing up number of like, number of retweet and then dividing this sum by the number of followers; this is further divided by the total number of tweets of the potential micro-influencer.

$$Engagement_i = \frac{\sum_t (Likes_{i,t} + Retweets_{i,t})}{TotalFollowers_i \times NumberofTweets_i} \quad (3)$$

This score measures the general dissemination ability of a micro-influencer. The higher the engagement the faster a message will spread inside and outside his community.

A subset of the gold standard dataset is presented in Table 1, where users have been anonymized with an auto incremental integer number in the first column. The next three columns report real scores computed with the rule previously defined in this section. The last column shows the label assigned to the user where 1 means micro-influencer and 0 means not micro-influencer. This section answer the RQ1 with the adoption of these three scores (Embeddedness, Interaction, and Engagement) and their relative thresholds for defining a gold standard in the micro-influencer field.

3.2 Mining personality traits and community-based values from text

When dealing with computational linguistics, we must know if micro-influencers use recurrent lexical expressions with respect to specific topics. We use a pre-trained word embeddings model from GloVe. GloVe has a fine-tuned model for Twitter.³

GloVe performs better with short text structures. Therefore, the word representation in terms of word embeddings is suited to our context. Tweets are characterized by short sentences containing many domain based tokens such as hashtags and handles.

We need a model that has been trained considering this context, otherwise we reduce the pre-training phase effectiveness in word embeddings definition.

²<https://www.grin.co>

³<http://nlp.stanford.edu/data/glove.twitter.27B.zip>

User	Embeddedness	Interaction	Engagement	Micro influencers
0	0.487231	0.0	0.000937	0
1	0.734960	0.0	0.000616	0
2	1.973224	0.833333	0.001889	1
3	0.808098	1.171428	0.063291	1
4	0.325655	0.0	0.000311	0

Table 1: Gold Standard example scores. Users, who have been anonymized with an auto incremental integer number, are listed in the first column. The next three columns report the actual scores computed on those users, as described in Section 3. Scores are rounded up to the sixth decimal place. The last column shows the label assigned to the user where 1 means micro-influencer and 0 means not micro-influencer.

We develop an algorithm to compute both BHV (Basic Human Values) and FFM (Five Factor Model) scores. Basic Human Values research [24] gives us example words for each human values: *selfdirection, stimulation, hedonism, achievement, power, security, conformity, tradition, benevolence, universalism*. We create a centroid per Basic Human Value averaging embeddings of words given as sample. In the case of benevolence, as an example, listed word are the following: *helpful, honest, forgiving, responsible, loyal, friendship, love, meaningful*. Each of these words has a 300 dimensions array representation in the GloVe pre-trained vocabulary. The first dimension of *helpful* is summed with the first dimension of *honest, forgiving, responsible, loyal, friendship, love* and *meaningful*; the sum is then divided by seven. This process is repeated for the next 299 dimensions. At the end, we obtain the 300 coordinates of the so computed centroid of the *benevolence* community value. After this phase we have 10 centroids, 1 per community based value.

After this configuration phase, one by one each word written by the micro-influencer candidate is parsed, as shown in Figure 2.

We compute the distances between this word and each community based reference centroid, then the word is assigned to the closest centroid in terms of euclidean distance. The number of words used in each Basic Human Value is multiplied by the inverse of the distance of the Basic Human Value centroid and the centroid obtained averaging spatial representation of all micro-influencer words assigned to that Basic Human Value. At the end, we have a score per BHV per micro-influencer.

$$\min(d(p, q_i)), i \in \text{SchwartzSampleCentroids} \quad (4)$$

Equation 4 computes euclidean distances between a word in analysis and Schwartz example centroids and select the minimum. The word is then assigned to the closer Basic Human Value cluster.

We compute final score for each community based values per user in Equation 5.

$$SS_i = nw_i * \left(\frac{1}{d(\text{avgEmb}_i, \text{Schwartz}_i)} \right) \quad (5)$$

We multiply the total number of words used in a community based value per the inverse of the euclidean distance between the example centroid and the averaged centroid of the words used in

that semantic cluster. At the end, we obtain 10 BHV scores per user. We show the first portion of the dataset containing the BHV scores of the analysed users in Table 3.

Five Factor Model [16] scores are retrieved adapting a method of Carducci et al. [5]. This method deals with the myPersonality dataset⁴ that contains FFM score coming from a psychological questionnaire paired with social media posts written by the same users who answered to the questionnaire. We train a Support Vector Machine algorithm for each of the FFM dimensions to retrieve the characteristics of writings related to psychological outputs. C and γ parameters in Support Vector Machine regression algorithm are selected with a grid search. This approach tests multiple values for C and γ and it selects with which of them the regression performs better. We can manipulate data samples considered by varying the C parameter: with small values we use almost every sample in the dataset, while high values are used to consider just the ones close to the margin of the hyperplane. Another important parameter is γ , which determine how flexible or rigid is the hyperplane, as it works indeed on the kernel. If the gamma value is too large then overfitting could occur. We have one regression model for each trait in Five Factor model. We use them to predict FFM personality traits of a micro-influencer starting from what he writes, and without having results from a psychological questionnaire.

We answer our RQ2 with the definition and the application of the two previous methods regarding the prediction of Five Factor Model and Basic Human Values scores. In fact, we extract personality traits and community based values of micro-influencers from their social media posts text.

4 EXPERIMENTAL EVALUATION

We use Twitter social media platform to retrieve user writings. In our setting, we select Twitter users having recently posted about trending topics. Users with less than 1k and more than 100k followers are removed. Out of this range a user is no longer considered a micro-influencer due to the volume of people in his audience [22]. For each user we download both his tweet text corpus and the list of ids of his followers.

We use the Natural Language Toolkit⁵ to perform data cleaning. In a first step, immediately after the downloading phase, we store tweets as a tsv (tab separated value) file removing all newlines and tabulations from the original text. In a second phase we apply the following procedures:

- stop-word removal: “Uhm, where is the leader? @johnsmith #officelife. :)” to “, where leader? @johnsmith #officelife. :)”
- punctuation removal: “,where leader? @johnsmith #officelife :)” to “where leader @johnsmith #officelife :)”
- emoticon removal: “where leader @johnsmith #officelife :)” to “where leader @johnsmith #officelife”
- handle and url removal: “where leader @johnsmith #officelife http://.../” to “where leader #officelife”

Hash tags (#) are preserved to highlight the topics in each processed tweet. Once cleaned, we tokenize the text at a space based level and we search each analysed token in the corresponding pre-trained and fine-tuned embeddings vocabulary.

⁴<https://sites.google.com/michalkosinski.com/mypersonality>

⁵<https://www.nltk.org>




Figure 2: BHV scores prediction pipeline. User tweets are cleaned and tokenized, each token has its own 300 dimensions array representations from GloVe embeddings. Each token array-shaped is assigned to the closest Basic Human Values reference centroid. Then, we sum up the number of words in each cluster and we multiply it by the inverse of the average position of these words. We obtain ten community based scores per user with this process.

We give tokenized and vectorized words as input to the FFM and BHV models described in Section 3, where we explain how to calculate them. This process creates as output five scores regarding FFM (Five Factor Model) and ten scores for BHV (Basic Human Values). The FFM and BHV scores are the inputs, while the micro-influencer labels previously computed are the expected outputs of three supervised classifiers, SVM, Random Forest, and CNN, and an ensemble model, XGBoost.

SVM is used in supervised learning to create an hyperplane that maximize the margin between two classes (in our case micro-influencer and not micro-influencer). As shown in Equation 6, Support Vector Machine tries to maximize the distance between the micro and not micro-influencer categories.

$$y = \sum_{i=1}^N (\alpha_i - \alpha_i^*) \cdot \langle \varphi(x_i), \varphi(x) \rangle + b \quad (6)$$

$$K(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\zeta^2}\right) \quad (7)$$

We choose a RBF (Radial Basis Function) as kernel (Equation 7), where x and x' represent two features of the model. The RBF kernel is a mathematical way to learn a non-linear classification rule that corresponds to a linear classification rule for the transformed data

points. The computation is done in a higher dimensional space to separate the classes and then it is projected to a lower dimension to see the transformed function.

In contrast, Random Forest is an ensemble method that operates by constructing a multitude of decision trees and outputting the class that is the mode of the classes or mean prediction of the individual trees. In our case the classes are two: micro-influencer or not. We use 15 as *max_feature*⁶ parameter deriving from 10 BHV scores and 5 FFM scores. We use a class-weight of 1:10 to deal with the unbalanced dataset. We give more weight to the micro-influencer class.

We test CNN (Convolutional Neural Network) as a classifier receiving as input BHV and FFM scores of users and as expected output the micro-influencer label. We build a sequential model in Keras with two layers and we use *adam* as adaptive learning rate optimization algorithm. We adopt *relu* as activation function. The loss function adopted is *cross-entropy*.

We use XGBoost,⁷ an optimized distributed gradient boosting library, to improve the classification performance. XGBoost uses an iterative approach to combine different models. It trains them consecutively, reducing errors made in previous steps.

$$\begin{aligned} \text{obj}^{(t)} &= \sum_{i=1}^n (y_i - (\hat{y}_i^{(t-1)} + f_t(x_i)))^2 + \sum_{i=1}^t \Omega(f_i) = \\ &= \sum_{i=1}^n [2(\hat{y}_i^{(t-1)} - y_i)f_t(x_i) + f_t(x_i)^2] + \Omega(f_t) + \text{constant} \quad (8) \end{aligned}$$

Equation 8 formalizes the XGBoost objective function. XGBoost makes splits up to the *max - depth* and starts pruning trees backward. This depth-first approach improves computational performance. It penalizes more complex models through both LASSO and Ridge regularization to prevent overfitting. XGBoost naturally admits sparse features for inputs by automatically learning best missing value depending on training loss and it handles different types of sparsity patterns in the data more efficiently.

We eventually choose SVM after looking at the results in terms of validation metrics, as shown in Table 5.

We use recall, precision and F1-score as validation metrics to understand how good is the classifier in detecting micro-influencers. Validation metrics are described in the following.

Recall highlights the vulnerability of the model in economical terms, because our tool finds micro-influencers among a huge amount of non micro-influencer users so, if the number of *false negatives* is high, it leads to few valid results. Recall needs to be as closer as possible to one so that number of false negative is low. Precision shows the presence of users who are considered micro-influencers but not acting as hoped. F1-score is the harmonic mean between recall and precision: it is useful to understand the compromise leading to not much effort in the manual cleaning of the micro-influencer wrongly predicted and, on the other side, not leaving out too many *true positive* results.

Validation metrics, in the binary case, compute how many candidates are well classified with respect to the expected output as

⁶The number of features to consider when looking for the best split.

⁷<https://xgboost.readthedocs.io/en/latest/>

described in the following. tp (true positive) represents the number of micro-influencers correctly classified. tn (true negative) represents the number of not micro-influencers correctly classified. fp (false positive) represents the number of not micro-influencers classified as micro-influencers. fn (false negative) represents the number of micro-influencers classified as not micro-influencers.

$$recall = \frac{tp}{tp + fn} \quad (9)$$

$$precision = \frac{tp}{tp + fp} \quad (10)$$

$$accuracy = \frac{tp + tn}{tp + tn + fp + fn} \quad (11)$$

We use stratified k-fold to compute validation metrics (recall, precision, f1-score).

Stratified k-fold is performed with Support Vector Machine, Random Forest Classifier, Convolutional Neural Network and XGBoost.

Stratified k-fold is a cross validation process where we divide the whole dataset in k subsets, we use $k - 1$ fold to train the classifier and the last one to predict outputs. We select k equals to 10 and at each iteration we store validation metrics. At the end of the process we average all validation metrics to obtain scores shown in Table 5.

Stratified k-fold is used because it allows us to maintain micro-influencer samples in every round of the fold. In fact, it emerges from the statistics shown in Table 2 that the classes are really unbalanced with few micro-influencers with respect to the total amount of users retrieved.

One of the main difficulties in this process is reflected in the statistics shown in Table 2. The proportion between micro-influencers and total user analysed is one over ten or less. As explained in Section 1, it is harder to find these users with respect to famous influencers. We worked with half million tweets per topic in our preliminary experiment, but just a small part of them are related to the topic used in the first query to Twitter, when we search micro-influencers talking about a trending topic.

These two considerations highlight the need of spending much time and effort in an initial phase when our dataset is small and when all new users searched are noisy. In a future scenario our system will need continuous tuning and updates while dealing with a collection of cases larger and more significant than the initial situation with few data available.

This initial effort is also mirrored in the results obtained by our classifier as shown in Table 5. This table presents the best performance in recall score, as desired in our particular context. Overall SVM performs better than the other three approaches. In the specific case of micro-influencers detection, recall measures how many users are considered not micro-influencers even if they are micro-influencers. If the metric is near to 1, then very few micro-influencers are not detected. Because micro-influencers are difficult to be found, it is important to detect all of them.

Higher results are obtained with *offgrid* topic because there the percentage of micro-influencer found is higher than the other topics. SVM recall is always higher than other classifiers recall. We notice CNN performs really bad in the case of *biodynamic* and *greenliving*. This is due to the shortage of micro-influencer cases. By looking at the precision score, it emerges how many not micro influencer are sometimes labeled as micro-influencers. This situation comes

Topic	Number of users	Micro influencers	Total tweets per topic
offgrid	146	10.96 %	407,957
plasticfree	190	6.32 %	560,655
biodynamic	70	8.57 %	201,010
greenliving	153	5.23 %	454,312
womenintech	238	3.78 %	644,772
sustainable	219	5.02 %	658,854

Table 2: Statistics about the defined micro-influencer gold standard dataset. Users are grouped by topic (hashtag searched). Second and third columns highlight the disproportion between total user retrieved and effective micro-influencer. Percentages are rounded up to the second decimal place. The last column represents the total number of tweets retrieved per topic summing all analysed user tweets.

from the *class_weight* parameter set in the stratified k-fold set to 1:10 because of the unbalanced presence of few positive cases with respect to the total amount of users analysed. Finally, XGBoost ensemble methods tends to favor precision with respect to recall and in our specific environment is not useful.

The classification of micro-influencers with a feature vector composed of personality traits and community based scores given as input to SVM classifier answer our RQ3.

5 DISCUSSION

Psychologists adopt Five Factor Model and Basic Human Values to compute personality traits and community values. We applied them to extract these scores from text and then we used them as features in micro-influencer detection. These models are selected as representative in the influence mechanism among social media users, the first on a single individual characteristics and the second on a broader community level. Our hypothesis is that a micro-influencer must possess special personality traits and he must share community values with the community he lives in.

Anyway, the automatic classification of personality traits introduces an initial error in the computation of the features, as highlighted in Carducci et al. [5]. The automatic classification of community values introduces further errors. To mitigate them, we plan to explore a validation dataset created with questionnaires.

Simultaneously, a shared definition of micro-influencers must come from a longer time-span experiment in which we monitor candidates and we score their performances dynamically. Nevertheless, a human validation process is not required to uniquely define a user as a micro-influencer because we start from the assumption that the rules of the gold standard already work as validation. These reasoning are useful to properly consider our results and to show that the research must focus on the improvement of these steps.

Word embeddings available in the released dataset⁸ are computed from two different pre-trained models: FastText⁹ for Five

⁸<https://doi.org/10.6084/m9.figshare.11309669>

⁹<https://fasttext.cc/docs/en/english-vectors.html>

User	selfdirection	stimulation	hedonism	achievement	power	security	conformity	tradition	benevolence	universalism
0	634.143	860.415	68.613	494.715	91.957	3290.369	527.879	28.087	836.087	476.888
1	520.099	594.802	72.731	347.464	88.334	3387.620	580.317	23.104	761.851	410.143
2	542.946	747.503	73.518	418.922	89.590	2754.447	553.618	39.676	878.727	585.977
3	255.612	256.765	4.789	156.130	14.810	955.365	281.866	9.873	238.837	166.703
4	796.353	569.985	90.697	202.510	145.414	3739.612	474.904	31.913	958.819	438.540

Table 3: Community based scores. These users are all potential micro-influencers remained after the first filter on the number of followers, as described in Figure 1. The anonymized identifier of the user is reported in the first column. The next ten columns show the final score obtained after user tweet corpus analysis following the procedures described in Section 3.

User	Openness	Conscientiousness	Extraversion	Agreeableness	Neuroticism
0	4.208	3.683	3.176	3.953	2.608
1	4.149	3.638	3.103	3.966	2.608
2	4.219	3.732	3.331	4.037	2.608
3	4.154	3.758	3.291	4.011	2.608
4	4.197	3.826	3.292	4.071	2.608

Table 4: Personality Traits scores. These users are all potential micro-influencers not discarded by the first filter on the number of followers, as illustrated in Figure 1. The anonymized identifier of the user is reported in the first column. The next five columns represent the final score obtained after the analysis of the user’s tweet corpus according to the procedures of Section 3.

Topic	SVM			Random Forest			CNN			XGBoost		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
offgrid	0.42	1.00	0.58	0.49	0.44	0.50	0.63	0.31	0.42	0.50	0.57	0.52
plasticfree	0.38	0.75	0.50	0.48	0.30	0.30	0.40	0.29	0.34	0.31	0.31	0.31
biodynamic	0.40	0.60	0.47	0.01	0.01	0.01	0.01	0.01	0.01	0.17	0.27	0.20
greenliving	0.26	0.50	0.32	0.50	0.13	0.20	0.01	0.01	0.01	0.23	0.13	0.17
womenintech	0.21	0.55	0.30	0.27	0.05	0.07	0.31	0.09	0.14	0.14	0.18	0.14
sustainable	0.23	0.47	0.30	0.48	0.11	0.14	0.6	0.27	0.38	0.13	0.16	0.13

Table 5: Experimental comparison of the considered classifiers. This table presents the best performance in recall score, as desired in our particular context. Overall SVM performs better than the other three approaches.

Factor Model and GloVe¹⁰ for Basic Human Values. These choices are motivated by experimental trails.

In terms of baselines, micro-influencer classification has no previous comparison references. Therefore, we compared different classifiers on standard quality metrics such as precision, recall and f1. Starting from these results, further studies should monitor improvements. Our work proposes the use of natural language processing techniques for deep social analysis in a not yet explored user niche of micro-influencers.

6 CONCLUSION

In this paper, we showed how it is possible to find micro-influencers and how to highlight their personality traits on individual (FFM) and community (BHV) level analyzing their writings. The whole

process can be reproduced to retrieve new data in different topics. Furthermore, changing pre-trained word embeddings and fine-tuning them in other fields of interest allows FFM and BHV scores extraction from different text sources.

Our best performing model presents high scores in recall, but we still need to improve its precision, maybe finding more examples of micro-influencers to better train the classifiers.

Future works in this field can adopt new computational linguistics algorithms and they can deal with new and shared definitions of micro-influencers. As described in Section 4, we removed emoticons during data cleaning. Therefore, additional studies are required to explore the loss of information introduced by this procedure and how much emoticons impact micro-influence.

Finally, further insights in this area are related to the analysis of audio and video features to highlight other traits useful for understanding the process of micro-influence.

¹⁰<http://nlp.stanford.edu/data/glove.6B.zip>

REFERENCES

- [1] Isabel Anger and Christian Kittl. 2011. Measuring influence on Twitter. In *Proceedings of the 11th International Conference on Knowledge Management and Knowledge Technologies*. ACM Press, New York, NY, USA, Article 31, 4 pages. <https://doi.org/10.1145/2024288.2024326>
- [2] Shlomo Argamon, Sushant Dhawle, Moshe Koppel, and James W. Pennebaker. 2005. Lexical Predictors Of Personality Type. In *In proceedings of the joint annual meeting of the Interface and the Classification society of North America*. In Proceedings of the 2005 Joint Annual Meeting of the Interface and the Classification Society of North America, Cincinnati, OH, USA, 16.
- [3] Eytan Bakshy, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. 2011. Everyone's an influencer: Quantifying influence on Twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining*. ACM Press, New York, NY, USA, 65–74. <https://doi.org/10.1145/1935826.1935845>
- [4] Carolina Bigonha, Thiago N. C. Cardoso, Mirella M. Moro, Marcos A. Gonçalves, and Virgilio A. F. Almeida. 2011. Sentiment-based influence detection on Twitter. *Journal of the Brazilian Computer Society* 18, 3 (2011), 169–183. <https://doi.org/10.1007/s13173-011-0051-5>
- [5] Giulio Carducci, Giuseppe Rizzo, Diego Monti, Enrico Palumbo, and Maurizio Morisio. 2018. TwitPersonality: Computing Personality Traits from Tweets Using Word Embeddings and Supervised Learning. *Information* 9, 5 (may 2018), 127. <https://doi.org/10.3390/info9050127>
- [6] Jilin Chen, Gary Hsieh, Jalal U. Mahmud, and Jeffrey Nichols. 2014. Understanding Individuals' Personal Values from Social Media Word Use. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. ACM, New York, NY, USA, 405–414. <https://doi.org/10.1145/2531602.2531608>
- [7] Easley David and Kleinber Jon. 2010. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, Cambridge.
- [8] Alpaslan Burak Eliacik and Nadia Erdogan. 2018. Influential User Weighted Sentiment Analysis on Topic Based Microblogging Community. *Expert Syst. Appl.* 92, C (Feb. 2018), 403–418. <https://doi.org/10.1016/j.eswa.2017.10.006>
- [9] Karen Freberg, Kristin Graham, Karen McGaughey, and Laura A. Freberg. 2011. Who are the social media influencers? A study of public perceptions of personality. *Public Relations Review* 37, 1 (2011), 90 – 92. <https://doi.org/10.1016/j.pubrev.2010.11.001>
- [10] Oliver P. John, Alois Angleitner, and Fritz Ostendorf. 1988. The lexical approach to personality: A historical review of trait taxonomic research. *European Journal of Personality* 2, 3 (1988), 171–203. <https://doi.org/10.1002/per.2410020302>
- [11] Christine Kiss and Martin Bichler. 2008. Identification of Influencers - Measuring Influence in Customer Networks. *Decis. Support Syst.* 46, 1 (Dec. 2008), 233–253. <https://doi.org/10.1016/j.dss.2008.06.007>
- [12] Upendra Kumar, Aishwarya N. Reganti, Tushar Maheshwari, Tanmoy Chakroborty, Björn Gambäck, and Amitava Das. 2018. Inducing Personalities and Values from Language Use in Social Network Communities. *Information Systems Frontiers* 20, 6 (Dec. 2018), 1219–1240. <https://doi.org/10.1007/s10796-017-9793-8>
- [13] Jure Leskovec, Lada A. Adamic, and Bernardo A. Huberman. 2007. The Dynamics of Viral Marketing. *ACM Trans. Web* 1, 1, Article 5 (May 2007), 39 pages. <https://doi.org/10.1145/1232722.1232727>
- [14] N. Majumder, S. Poria, A. Gelbukh, and E. Cambria. 2017. Deep Learning-Based Document Modeling for Personality Detection from Text. *IEEE Intelligent Systems* 32, 2 (Mar 2017), 74–79. <https://doi.org/10.1109/MIS.2017.23>
- [15] Amit Mandelbaum and Adi Shalev. 2016. Word Embeddings and Their Use In Sentence Classification Tasks. *CoRR abs/1610.08229* (2016), 16. arXiv:1610.08229 <http://arxiv.org/abs/1610.08229>
- [16] Robert R. McCrae and Paul T. Costa. 1987. Validation of the five-factor model of personality across instruments and observers. *Journal of Personality and Social Psychology* 52, 1 (1987), 81–90. <https://doi.org/10.1037/0022-3514.52.1.81>
- [17] Robert R. McCrae and Paul T. Costa. 2008. Empirical and Theoretical Status of the Five-Factor Model of Personality Traits. In *The SAGE Handbook of Personality Theory and Assessment: Volume 1 – Personality Theories and Models*. SAGE Publications Ltd, Los Angeles, CA, USA, 273–294. <https://doi.org/10.4135/9781849200462.n13>
- [18] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *arXiv e-prints* 1, 1, Article arXiv:1301.3781 (Jan 2013), 12 pages. arXiv:cs.CL/1301.3781
- [19] James W. Pennebaker and Laura A. King. 1999. Linguistic styles: Language use as an individual difference. *Journal of Personality and Social Psychology* 77, 6 (1999), 1296–1312. <https://doi.org/10.1037/0022-3514.77.6.1296>
- [20] James W. Pennebaker, Matthias R. Mehl, and Kate G. Niederhoffer. 2003. Psychological Aspects of Natural Language Use: Our Words, Our Selves. *Annual Review of Psychology* 54, 1 (2003), 547–577. <https://doi.org/10.1146/annurev.psych.54.101601.145041>
- [21] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Stroudsburg, PA, USA, 1532–1543. <https://doi.org/10.3115/v1/d14-1162>
- [22] Monika Ewa Rakoczy, Amel Bouzeghoub, Alda Lopes Gancarski, and Katarzyna Wegrzyn-Wolska. 2018. In the Search of Quality Influence on a Small Scale – Micro-influencers Discovery. In *On the Move to Meaningful Internet Systems. OTM 2018 Conferences*, Hervé Panetto, Christophe Debruyne, Henderik A. Proper, Claudio Agostino Ardagna, Dumitru Roman, and Robert Meersman (Eds.). Springer International Publishing, Cham, 138–153.
- [23] Roy Schwartz, Roi Reichart, and Ari Rappoport. 2015. Symmetric Pattern Based Word Embeddings for Improved Word Similarity Prediction. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*. Association for Computational Linguistics, Beijing, China, 258–267. <https://doi.org/10.18653/v1/K15-1026>
- [24] Shalom H. Schwartz. 2012. An Overview of the Schwartz Theory of Basic Values. *Online Readings in Psychology and Culture* 2, 1, Article 11 (2012), 20 pages. <https://doi.org/10.9707/2307-0919.1116>
- [25] M. Su, C. Wu, and Y. Zheng. 2016. Exploiting Turn-Taking Temporal Evolution for Personality Trait Perception in Dyadic Conversations. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 24, 4 (April 2016), 733–744. <https://doi.org/10.1109/TASLP.2016.2531286>
- [26] Max Weisbuch, Zorana Ivcevic, and Nalini Ambady. 2009. On being liked on the web and in the “real world”: Consistency in first impressions across personal webpages and spontaneous behavior. *Journal of Experimental Social Psychology* 45, 3 (2009), 573–576. <https://doi.org/10.1016/j.jesp.2008.12.009>