Detecting Influential Users in Social Networks: Analysing Graph-Based and Linguistic Perspectives

Kévin Deturck^{1,2}, Namrata Patel^{1,3}, Pierre-Alain Avouac¹, Cédric Lopez^{1,4}, Damien Nouvel², Ioannis Partalas^{1,5}, Frédérique Segond^{1,2}

¹ Viseo, Grenoble, France

`first_name.last_name@viseo.com`

² INaLCO-ERTIM, Paris, France

`first_name.last_name@inalco.fr`

³ Université Paul Valéry, Montpellier, France

`first_name.last_name@univ-montp3.fr`

⁴ Emvista, Montpellier, France

`first_name.last_name@emvista.com`

⁵ Expedia, Lausanne, Suisse

`first_name.last_name@gmail.com`

Abstract. There has been increasing interest in the artificial intelligence community for influencer detection in recent years for its utility in singling out pertinent users within a large network of social media users. This could be useful, for example in commercial campaigns, to promote a product or a brand to a relevant target set of users. This task is performed either by analysing the graph-based representation of user interactions in a social network or by measuring the impact of the linguistic content of user messages in online discussions. We performed independent studies for each of these methods in the present paper with a hybridisation perspective. In the first study, we extract structural information to highlight influence among interaction networks. In the second, we identify linguistic features of influential behaviours. We then compute a score of user influence using centrality measures with the structural information for the former and a machine learning approach based on the relevant linguistic features for the latter.

Keywords: influence, social media, linguistics, centrality

1 Introduction

An influencer can be characterised as a person that has the power to affect people, actions or events. In recent years, There has been increasing interest in the artificial intelligence (AI) community for influencer detection in recent years for its utility in singling out pertinent users within a large network of social media users. Such information is crucial in many research studies such as in sociology and information management domains.

Additionally, with the frenetic growth of available data in social media, being able to analyse and detect influential users becomes crucial as they are susceptible to express their ideas with a greater impact than other individuals. This could be useful, for

example in commercial campaigns, to promote a product or a brand to a relevant target set of users and maximise their spread [1].

Influencer detection is usually performed by analysing a graph-based representation of user interactions in a social network. In this context, studies using graph theory leverage the structural information in these graphs to identify the most important nodes in a network [2] [3].

Following another line of thought, a recent development in the task has been to analyse the textual content of the messages posted by the users to identify characteristics of communication for influence detection [4][5].

We further explore both aspects in this paper with two independent studies, each generating the following specific contexts, based on their respective requirements: (1) a social network featuring various interaction types for graph analysis and (2) a forum that provides written posts as the only means of expression, for linguistic analysis. These two studies respectively allow us to (1) observe the interaction types and users' positions in graphs that highlight influential users and (2) address the influence of the linguistic content of user's messages.

Regarding graph analysis, we compare various interaction types in their capacity to denote influence between users. We build the graph from selected interactions to get their structure in a formal way. From this structure, we analyse users' positions and determine those that reflect influence among interactions. We use centrality measures considering a central position in interaction networks as an indicator of influence.

In developing our linguistic approach, we compare linguistic criteria (such as a user's argumentation, agreement / disagreement between users) with classical numerical criteria (number of answers, message size, number of relations, *etc.*). The former is extracted using a symbolic approach based on a set of linguistic rules while the latter is extracted from available metadata on the messages. This information is then integrated into an automatic learning system. The resulting system is thus a "doubly hybrid" system, since it is based on symbolic and statistical methods on the one hand, and information structure and textual content on the other. To facilitate the interpretation of results and to better represent the different aspects addressed in our approach, we complete our system with an interface for knowledge visualisation.

The rest of the article is organised as follows. In section 2 we provide a general overview of existing techniques in influence detection within social media, situating our approaches in their scientific context. We then describe our approaches in section 3 and follow it up with their respective evaluations in section 4.

2 Related Work

With the growing number of users interacting through social media and the increasing amounts of available data for research, social media information has become a considerable source for user behaviour analysis. For influence detection, in particular, social media contains a significant amount of information which can be exploited at the level of (1) user interactions and (2) message content. This has led to two main lines of work.

2.1 Influencer Detection by Social Network Analysis (SNA)

The first axis of work in this domain is primarily based on graph theory, where influential behaviour is computed by analysing the structural information of user interactions contained within a given sample of a social network. SNA meets graph theory as user interactions are formally represented in a graph with users as vertices and interactions as edges. Among the different types of social media, including forums, blogs and social networks, the latter are commonly studied for graph analysis as it offers more diverse user interactions. For example, Twitter provides “follow”, “like” and “reply” while forums or blogs only provide “reply” and “quote” as interactions.

[6] introduces the idea that there is a connection between an individual’s central position among the interactions of a group and their influence within the same group. [7] formalises the centrality intuition and proposes a first set of graph centrality measures, including *degree*, *proximity* and *betweenness* measures. Centrality measures compute a centrality score for each vertex in a graph according to a specific edge configuration. For example, the *degree* measure uses the number of direct edges from a vertex while *proximity* measures the average path length (number of edges) from one vertex to others. The centrality score of a vertex indicates its influence value in a graph. With the growing popularity of web pages, a new centrality measure was defined, PageRank [3], to take into account hyperlinks between pages. Applying different centrality measures on social graphs which represent different interaction types provides various aspects of influence measurement. [8] applies a proximity measure on scientific citation networks to measure the influence of researchers. [9] highlights the importance of the community to analyse centrality. Authors claim that users of a social network tend to communicate with users of a same group, for example regarding a specific topic. They detect communities in a social network sample before computing users’ influence value in each community. [10] fine-tunes the selection of interactions distinguishing user-level from content-level interactions in a multilayer adaptation of PageRank.

Some works defined influencers according to their capacity to motivate individuals towards an action or a message, thus detecting them by analysing the propagation of interactions. In a graph, an influential vertex could replicate its behaviour in a fast and deep way. This vision relates to the key problem of influence maximisation through a network. [11] predicts the interaction dynamics from a graph node to estimate its influence. [12] analyses information spread to detect influencers. [13] associates influencers with users who are able to maximise their opinion through a network.

2.2 Influencer Detection by Linguistic Analysis

The second axis of work delves into the semantic aspects of the user messages, identifying influential behaviour through linguistic markers. Recent research in this category aims at identifying characteristics of influential behaviour through linguistic markers present within the messages. [14] focuses on the opinions expressed in messages to follow influential trends. [15] and [5] describe several behavioural features such as persuasion, agreement/disagreement, dialog patterns and sentiments which characterise

influence and propose a machine learning approach to detect influential users. [16] identifies influencers by a specific language use including emotion lexicon and personal pronouns to establish a proximity with their audience, thus facilitating the message transmission.

As we mentioned the importance of community in SNA state-of-the-art, linguistic analysis helps with finding communities regarding a topic. [17] drives influence detection according to topic detection in users' messages.

This line of work provides promising results in influence detection, given the depth of scrutiny involved in the analysis of influential behaviour that we can relate to particular effects for the audience like opinion change.

2.3 Influencer Detection by Hybridisation

The challenge of combining both axes of research is relatively less explored. [18] biases PageRank towards certain users according to a specific topic. More recently, [19] proposed a supervised random walk approach towards topic-sensitive influential nodes. As can be seen, the message content is exploited here only in terms of the topic. Taking into account the second line of research, this challenge can be addressed by focussing on the semantic aspects of the message content.

In the aim of combining these two lines of research, as a first step, we independently explore each in their own context. These can then be capitalised upon with a hybridisation perspective. Focusing on the semantic aspects of message content, we develop a linguistic rule-based reasoning engine to identify linguistic markers for influential behaviour in a corpus of forum discussions. As for the SNA approach, we propose an exploration of some centrality measures applied on a manually annotated Twitter¹ dataset in order to determine those centrality measures and Twitter interactions that are relevant to influence detection.

Performing these studies independently allows us to analyse them respectively in favourable conditions, as each of them has specific requirements that are not compatible with the other. For example, graph-based studies require complex interactions which forums cannot provide and linguistic studies require substantial textual content which Twitter cannot provide.

3 Methodology

We propose an exploration of the two main ways to detect influencers. (1) Linguistic analysis and (2) centrality computation are tested separately on specific datasets that conform to the respective requirements of the two types of analysis: (1) a forum that allows long messages for linguistic analysis and (2) Twitter featuring different interaction types for SNA. Both experiments can be broken down into the four following phases:

1. Corpus construction from social-media source

¹ <http://www.twitter.com>

2. Selection of linguistic and structural features
3. System design
4. Visualisation

We now describe each work in further details.

3.1 Corpus Construction

Forum Dataset

The data used to elaborate the influencer detection algorithm comes from an English forum in the domain of cosmetics² which contains different discussions about makeup products, beauty tips, *etc.* We have scraped more than 5,000 threads from this forum and randomly divided the corpus in three different groups. The first group, *RuleDevelopment*, consists of 1000 threads reserved for analysis and to develop linguistic rules; the second group, *TrainingSet*, also consists of 1,000 threads and serves as training data for the machine learning module; the third group, *TestSet*, consists of the rest of the threads (3,000) and is used to evaluate our approach. Each of the 18,085 messages within the second group (dedicated to training the machine learning model and 1027 messages within the third group was manually annotated to reflect a boolean value per message: whether or not the message is influential.

We consider as influential a message that contains specific linguistic features reflecting influential behaviour (further described in section 3.2). We therefore defined an annotation guide specifying how an annotator may recognize these features. For example, the following message shows that its author has been *argumentative*. “What is the look for this season? Is the dewy face in or a matte one? I see some stars with make up where their face is dewy and *it looks nice, but I can't stand doing that. I have to make my face matte or otherwise I feel shiny and oily.*”

Twitter Dataset

As a ground-truth, we chose a dataset made for RepLab 2014 [20]. This lab included a task that consisted in ranking Twitter users according to their real-world influence. The dataset contains more than 7,000 accounts split according to their domain: bank, automotive and others. Accounts were binary annotated depending on whether they were real-world influencers or not by the online reputation experts Llorente & Cuenca³. The dataset contains in average 1/3 of influencers.

² We dissimulate the name of the forum for reasons of confidentiality.

³ www.llorentecuenca.com/en/

3.2 Features

In section 1, we defined influence as a power, that leads us to characterise influencers according to the resources and effects of this power. We look for these two aspects analysing interaction content and structure.

Features from the Forum Dataset

During this phase, the corpus is analysed to identify criteria related to influential behaviour, as cited in the section above and described in Table 1. We distinguish between "linguistic" and "non-linguistic" criteria to separate the linguistic information from the structural one. The former is extracted on the basis of a set of linguistic rules. The latter is computed using count functions or by determining a boolean value.

Table 1. Description of the features extracted to be used in the machine learning model.

Category	Type of the features	Features	Output
Non-linguistic features	Position of the post in a thread	isFirstPost?	boolean
Non-linguistic features	Position of the post in a thread	isSecondPost?	boolean
Non-linguistic features	Position of the post in a thread	isPenultimateost?	boolean
Non-linguistic features	Position of the post in a thread	isLatestPost?	boolean
Non-linguistic features	Quantitative information	sizeOfMessages	integer
Non-linguistic features	Dates	RegistrationDate	date
Non-linguistic features	Location	Location of the user	string
Linguistic features	Writing style	Elongation	boolean
Linguistic features	Writing style	Uppercase	boolean
Linguistic features	Writing style	Exclamation	boolean
Linguistic features	Writing style	Interrogation	boolean
Linguistic features	Writing style	Advising	boolean
Linguistic features	Argumentation	Nb of premises	integer
Linguistic features	Argumentation	conclusion?	boolean
Linguistic features	Argumentation	ArgumentInFirstSentence	boolean
Linguistic features	Agreement/Disagreement	Agreement	boolean
Linguistic features	Agreement/Disagreement	Disagreement	boolean

To extract the linguistic features, we develop a separate module for each type of feature. We have 3 modules: (1) Writing style, (2) Argumentation and (3) Agreement/Disagreement. Each of these modules consists of the linguistic rules specific to the corresponding linguistic feature, developed by analysing the portion of the corpus kept aside for this purpose (*RuleDevelopment*). All the linguistic rules are based on a morphosyntactic analysis performed by the Eloquant Semantic Solutions⁴ parser. We now detail each of these linguistic modules.

Argumentation

To detect instances of argumentation within the messages, we base ourselves on the study described in [20]. An argument is defined as a set of propositions, each of them being a premise, with at most one being a conclusion.

Thus, we focus on the identification of messages that potentially contain premises and/or conclusions. For instance, "This product is not reliable and very expensive!" is a premise, and "Then I can't recommend buying it!" is a conclusion.

Writing Style

To extract features corresponding to "writing style", we exploit the way in which authors express their opinions. We detect four indicators of writing style.

- *Elongation*, e.g. "greeeeeeeeat"
- *Uppercase*, e.g. "I LOVE this product"
- *Exclamation/Interrogation*, e.g. "You should try it!!!"
- *Advising*, e.g. "You can buy this product"

Agreement/Disagreement

We develop the Agreement/Disagreement module on the basis of the following question: Does the author agree/disagree with previous author? For instance, in the following sentence: "I'm not going the same way as Mary", the system might be able to detect a disagreement.

All the rules developed for the different linguistic modules follow the same general pattern and are adapted according to the linguistic feature to be extracted. This pattern is described as:

1. *Construction of lexicons based on the state-of-the-art i.e.* detection of premises : "as shown by", "is implied by", "on the supposition that", "may be deduced from", ... ; Detection of conclusions: "concludes", "proves", "entails", "lead me to believe that", "bear out the point that", "it must be that", ...

⁴ <https://www.eloquant.com/en/semantic/explore-automatic-semantic-analysis>

2. *Morphosyntactic analysis with Eloquent parser* : we use the lemma and the form in order to take into account variations such as “is implied by”, “was implied by”, ...
3. *Application of rules* destined to detect whether a phrase from one of the lexicons (and all its variations) appears in a given message.

The messages are thus automatically annotated according to the different detected features. These then serve as input for the machine learning model that computes an influence score per message.

Features from the Twitter Dataset

As influence needs interactions (unidirectional actions from a user to another user), we focus on selecting Twitter interactions as features. We present Twitter interaction types in Table 2 excluding private message as, by definition, it is not accessible. We put in bold the types we finally selected. The principle of this selection was to work on interaction types that denote an influence on issuers from receivers. For each type, we analyse the engagement it implies for issuers towards receivers. *Follow* is the only interaction type that directly denotes an engagement at user-level. As influencers are basically users, we retain it. We decide to compare this user-level interaction type with one at content-level. Of the other three types, *answer* is the one that provides the least semantics by itself, requiring linguistic analysis; as this is not the point of the SNA exploration, we eliminated it. *Retweet* and *Like* have similar constitutions (cf. Table 2) and both indicate an interest from issuer for a content produced by receiver. However, we noticed that *retweet* is a stronger engagement as it affects a larger audience (cf. Table 2), we therefore retain it as a second feature.

Table 2. Twitter public interactions.

Nom	Source Type	Target Type	Main Effect	Audience
<i>Retweet</i>	User	Tweet	Related Tweet	As for source
<i>Like</i>	User	Tweet	Bookmark	As for source
<i>Answer</i>	Tweet	Tweet	Related Tweet	Twitter
<i>Following</i>	User	User	Subscription	As for source

3.3 System Design

System Design for Linguistic Analysis: Machine Learning Model Generation

We presented in section 3.2 the identification of linguistic and non-linguistic features to detect an influential behaviour in the forum dataset. During this phase, each message

is described in terms of the features it contains. The entire dataset is therefore represented as a matrix: each line represents a message and each column represents a feature that it contains. Feature values for a given message are filled in on the basis of the annotations present within it.

This feature matrix is fed to the machine learning model in order to compute the final influence score per message. We chose to employ Random Forests (RF) as they are proven to be robust and state-of-the-art methods across several applications. Essentially, a random forest algorithm creates multiple decision trees by learning simple rules. A decision of membership is made according to prediction probabilities. Fig. 1 presents a simple decision tree where the nodes in rectangles represent the leaf-level nodes and the prediction probability for a message to be influential is visually represented in burgundy.

Fig. 1. A simple decision tree computed by our model.

The procedure described above is applied to each message in the corpus. Therefore, as output at this stage, we have an influence score per message which represents the probability of responding positively to the question "Is this message influential?". These influence scores are then aggregated to produce a final influence score per author. This aggregation is done by exploiting the structural information present in the network of user interactions (authors).

Let $U = \{u_1, u_2, \dots, u_n\}$ be the set of users in a social network and $S_u = \{s_1, s_2, \dots, s_{K_u}\}$ be the set of scores for each post of user u , where K_u = number of messages posted by u . Then we define the following normalised aggregated value as the final influence score for each user:

$$Inf(u) = \frac{\frac{1}{K_u} \sum_{i=1}^{K_u} s_i}{\max_{u'} \sum_{j=1}^{K_{u'}} s_j \frac{1}{K_{u'}}} \quad (1)$$

Formalisation for Centrality Computation: Graph Construction and Centralities

As we detect influencers as particular Twitter users, we only build user graphs, even for *retweet* which is at content-level. In a graph, we represent both interactions described in section 3.2 as orientated edges between users as nodes (from issuer to receiver). We build one graph per interaction type.

We distinguish two types of centrality: local, that only takes into account direct links from a node, and global, that uses the whole graph. We select six centrality measures from these two types because of their complementarity and popularity. Both constitute different influence models that we detail below.

- **Incoming degree**: directly uses the number of incoming links for a given node. It is a direct polarization towards a user [7].
- **Betweenness and Proximity**: compute the shortest paths between all nodes in a graph to respectively get how much a node is present on it and how distant a node is on average from other ones [7]. We did not get results for these measures because we obtained disconnected graphs on which we cannot compute shortest paths between all nodes. Nonetheless, we mention them because they still are part of our experiments.
- **Hits** : is more complex as it computes two mutual scores: authority and hub. Authorities receive interactions in particular from hubs whereas hubs connect to a lot of authorities [21].
- **PageRank** : computes the centrality of a node taking into account not only the edge configuration but also the weight of edges and nodes with a uniform probability to “jump” between nodes [3].
- **LeaderRank** : is a PageRank adaptation for social networks setting the “jump” probability from a uniform value to one that decreases for each node with its number of outgoing links [22].

We experiment by applying the six centrality measures on both interaction types to find those combinations that best reflect the ground-truth. We independently apply each centrality measure on graphs. Centrality measures browse all nodes in a graph and compute a centrality score for each node according to their respective centrality formalisation. We strictly use centrality scores as influence values for the corresponding users to finally rank them by influence.

3.4 Visualisation

For a comprehensive visualisation of influential users, we develop a knowledge visualisation module which describes (1) a graphical representation of the users ranked by influence, (2) the user interactions present in a given thread (3) the message threads annotated to highlight the detected linguistic features. The module is web-based to allow for straightforward accessibility.

Based on the linguistic analysis performed by the system under the hood, the module presents a special feature that offers the possibility of filtering of social users according to *key-terms* or *topics of interest*. This gives a fine-grained view on the set of

influencers. The module also provides a global view of the detected influencers using different visualisation methods such as a bar chart or a bubble chart.

Fig. 2 presents a screen of the visualization module where the top 20 users are ranked in a bar chart according to their score of influence. Fig. 3 presents the interaction among users in a graph for a certain discussion. Users with higher score of influence are represented with bigger circles.

Apart from visualising the set of users, the module also offers a view of the message threads analysed to compute an influence score. The features used as input for the machine learning model are highlighted within the sentences to offer a contextual explanation of message content that is relevant to computing the final influence score.

The visualization module was actually made for a demonstration of the Soma project⁵ which does not feature centrality computation to compute scores of influence. Pursuing our hybridisation perspective, we intend to extend the visualization module so that one can select the features used to compute users' influence score including their centrality score. Nonetheless, we save Twitter data in a Neo4j⁶ database that offers a graph browser to visualize the information one requests. Fig. 4 and 5 respectively present screenshots of follow and retweet graphs with yellow circles as user nodes and blue circles as tweet nodes. We can observe that these two graphs are more complex than the conversation one presented in Fig. 3 thus analysing their structure with centrality measures seems particularly relevant.

Fig. 2. Influencer Ranking Visualisation.

⁵ The linguistic study was a part of the SOMA Eurostars project (SOMA Eurostars program 9292/12/19892, <http://www.somaproject.eu/>) which concerns the enhancement of customer relationship management systems with social media analysis capabilities.

⁶ <https://neo4j.com>

Fig. 3. The interaction graph of users throughout a discussion.

Fig. 4. The follow graph between Twitter users.

Fig. 5. The retweet graph between two Twitter users.

4 Evaluation

Our primary aim in evaluating our approach is to assess the relevance of including a linguistic analysis of message content in the detection of influencers. We therefore perform experiments that compare the performance of our developed system with and without the linguistic analysis. We also rank the influence criteria used as features in the ML model by order of importance to identify how linguistic features fare compared to the traditionally used numerical ones.

We evaluate the different combinations of centrality measures and interactions types by comparing user rankings they produce against the RepLab binary reference as an Information Retrieval (IR) task: best rankings must return the most of influencer users in higher positions. We now describe the experimental setup and the obtained results.

4.1 Experimental Setup

Linguistic Approach

To evaluate the proposed approach we used the corpus described in section 3.1. During evaluation we omitted *RuleDevelopment*, the portion reserved to develop linguistic rules, to avoid any resulting bias. We first trained the ML model using the manually annotated TrainingSet, then tested the model on TestSet. Our selected ML model, Random Forests, also allowed the extraction of feature importance. This is particularly

useful to evaluate the pertinence of the linguistic features used in computing an influence score.

To train the machine learning model we used RF for which we performed a random search coupled with 5-fold cross validation to tune its parameters: (1) number of trees tuned to the interval $\in [50, 500]$, (2) the depth $\in [2, 10]$ and (3) the information criterion $\in \{entropy, gini\}$. We trained two versions of the model, with and without linguistic features in order to assess their relevance. The two models were optimized for ROC-AUC which is a measure of the probability that a positive instance will be ranked higher than a negative one.

Centrality Approach

Limited by Twitter API usage rates⁷ for extracting interactions, we decided to sample the RepLab dataset at 50 users maintaining the original 1/3 influencer proportion for comparability purposes with the lab participant systems and filtering by bank domain to get a community-like representation.

Building one graph for each interaction type allows us to compare centrality measures among them. We extracted all *follow* relationships regarding our sample and represented it in a *follow graph*. Considering *retweet*, we selected the last ten tweets from each user’s timeline and extracted the retweets in the limit of 100 imposed by Twitter API. To build a user graph while *retweet* is a content-level interaction, we represented as nodes the retweet authors instead of the retweets itself, assigning a unique edge in the case of multiple retweets between the same two authors to favour diversity to quantity.

We present some statistics on *follow* and *retweet* graphs respectively in Table 3 and 4. *Node count* includes the initial 50 root users, the rest being new users who interacted with root users. *Edge count* provides the number of unique interactions (we chose uniqueness for *retweet* and *follow* is unique by definition) between users. *Density* measures how much *potential graph space*, according to the number of nodes, interactions fill (cf. equation 2 with *Density*, the density for orientated graphs, *edgeCount*, the number of edges and *nodeCount*, the number of nodes). A first observation is the low density of both graphs that reflects a low interaction level in our sample.

$$Density = \frac{edgeCount}{nodeCount * (nodeCount - 1) / 2} \quad (2)$$

Table 3. *Follow graph* statistics.

<i>Number of Nodes</i>	5.067,480
<i>Number of Edges</i>	5,149,491
<i>Density</i>	10^{-7}

⁷ <https://developer.twitter.com/en/docs/basics/rate-limits.html>

Table 4. Retweet graph statistics.

<i>Number of Nodes</i>	2099
<i>Number of Edges</i>	2051
<i>Density</i>	10^{-4}

By definition, centrality measures do not give a direct answer to the question: *Who is an influencer*. Therefore, we preferred to rank users according to the scores rather than setting an empirical threshold which could skew the real answer to this question. To better get the contribution of centrality measures to influencer detection, we built a random user ranking as a baseline. As reference is based on a binary annotation instead of a ranking, we still need a specific metric to evaluate the final rankings. We thus selected Mean Average Precision (MAP) which was used for the same purpose in RepLab 2014. It is based on the IR principle which states that the most relevant results (here, the influencers) should appear in the highest positions of a ranking returned by a query system (here, requesting for influencers among some Twitter users). We compute MAP according to equation 2 with N , the total user count, n , the count of influencers correctly retrieved, $p(i)$, the precision at rank i (considering only the first i users), $R(i) = 1$ if the user at rank i is actually an influencer otherwise $R(i) = 0$.

$$MAP = \frac{1}{n} \sum_{i=1}^N p(i) R(i) \quad (3)$$

4.2 Results

Linguistic Approach

The results obtained indicate two ways of evaluating the pertinence of the linguistic features: (1) ROC-AUC curves comparing the system with and without the use of linguistic features and (2) ranking the features by order of importance to locate the position of linguistic features.

Fig. 5 presents the ROC-AUC curves for both systems. We note that allowing a false positive rate of 30% the system with the linguistic features can reach a true positive rate of 82%.

Fig. 6 presents the feature importance as extracted from the RF model. We note that the most important feature is the size of the post which naturally reflects the fact that linguistic phenomena such as argumentation or elongation are employed more often by users in longer messages. Interestingly, the use of premises (argumentation feature) as well as elongation (writing style feature) are shown to be important features.

We can thus see that the chart reveals that our selected linguistic features (argumentation and advising) find their place between two of the most often used non-linguistic features. This traces the path for the next steps in our research.

Fig. 6. ROC-AUC curve for both models.

Fig. 7. Feature importance ordered by decreasing significance in the RF model.

Centrality Approach

We present in Table 5 the results of centrality measures on the *follow graph*. The statistics of the *follow graph* (cf. section 4.1, Table 3), highlighting the low graph density, indicate that users do not follow each other a lot among the sample. PageRank and LeaderRank, which are global measures, take advantage from deep information about interactions through the graph but because of low density, they do not much improve results compared to Incoming degree, a local and simpler measure.

Hits provides the best result distinguishing influencers as *authorities*, thus showing relevance of the specific authority/hub relay for influencer detection and characterising influencers as users who are particularly followed by *hubs*.

Table 5. Results of centrality measures on *follow graph*.

Centrality	Baseline	Incoming degree	PageRank	LeaderRank	Hits
MAP (%)	38,67	43,49	44,25	44,53	51,68

We present in Table 6 the results of centrality measures on the *retweet graph*. The low graph density issue seems more difficult here because all global measures obtained the same results as the local one, *Incoming degree*. We explain it by the fact that information volume is lower for *retweet* thus it could make it even more difficult to be discriminative.

Hits does not get therefore the best result on the *retweet graph* as on the *follow graph*. This can be explained, as we previously indicated, by graph statistics rather than interaction types.

Table 6. Results of centrality measures on *retweet graph*.

Centrality	Baseline	Incoming degree	PageRank	LeaderRank	Hits
MAP (%)	38,67	40,91	40,91	40,91	40,91

5 Conclusion

In this paper, we have presented two studies that browse the two main aspects of the influencers detection state-of-the-art and explore (1) the influence of linguistic features in forum conversations and (2) the importance of Twitter interaction types regarding centrality measures to find out particular user's positions reflecting their influence.

We designed a hybrid approach for the detection of influencers based on symbolic and statistical methods on the one hand and on the structure and textual content of the networks on the other hand. Our aim has been to address the significance of exploiting linguistic criteria (such as a user's argumentation, agreement/disagreement between users) for influence detection with respect to the traditionally used numerical criteria (number of responses, message size, number of relations, etc. Our results confirm the relevance of the former in the detection of influence: the linguistic features pertaining to argumentation and writing style (in particular elongation) appear to be among the most relevant criteria.

We combined a selection of classical and complementary centrality measures on the one hand and Twitter interaction types we related to influence on the other hand. The

goal has been to identify the combinations that help the most in identifying influential users in the interaction graphs we built. SNA shows a low interaction density issue while we selected root users from the same domain (bank). We shall increase the graph density to make the centrality measures more relevant. We could either try to « sum up » the graph by removing nodes with particularly low degree that do not give important information or build graphs from proper community detection.

The two independent studies allowed us to analyse two distinct types of approach that we intend to conjugate. Regarding both studies we have just presented, we consider various solutions to build a hybrid system based simultaneously on linguistic analysis and SNA.

A first solution would be to adapt the model learnt on forum to apply it on Twitter. We shall select linguistic features according to current results taking into consideration Twitter specificities. We shall evaluate it on tweets from the same user sample we used in this paper for comparability purposes. We shall also compute a hybrid influence score combining the aggregated score by user with (1) a centrality score obtained with the best structural modelling observed in this paper or (2) several centrality scores we shall weight according to their current respective result.

A second solution, more difficult to implement and that we shall also apply on the RepLab dataset, consists in enriching graph representation with linguistic information. For example, we shall represent the agreement between a user *A* and a user *B* by an edge from a node *A* to a node *B*. It is also possible to weight edges according to the importance we give to the information types they represent, including both raw and linguistic-based information. We shall finally apply centrality measures on the resulting multilayer graph.

The third and last solution is a combination at a larger scale as it includes a previous work on a dataset from the Change My View forum [23]. In this work, we have presented our approach to opinion change detection with the goal of identifying influencers by their effect, thus considering opinion change as an influencer’s effect. We shall apply the linguistic model and centrality measures we have presented in this paper combining them with the opinion change detection module to detect influencers in the Change My View forum.

References

1. Richardson, M., Domingos, P.: Mining knowledge-sharing sites for viral marketing. In: Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 61-70. ACM (2002).
2. Bonacich, P.: Power and Centrality: A Family of Measures. *American Journal of Sociology* 92(5), 1170-1182 (1987).
3. Page, L., Brin, S., Motwani R., Winograd, T.: The pagerank citation ranking: Bringing order to the web. In: Technical Report 1999-66 Stanford InfoLab, Stanford (1999).
4. Kien-Weng Tan, L., Na, J. C., Theng, Y.L.: Influence detection between blog posts through blog features, content analysis, and community identity. *Online Information Review* 35(3), pp. 425-442 (2011).
5. Rosenthal S.: Detecting Influencers in Social Media Discussions. PhD thesis, Columbia university, Columbia (2015).
6. Bavelas, A.: A mathematical model for group structures. *Applied anthropology* 7(3), 16-30 (1948).
7. Freeman, L. C.: Centrality in social networks conceptual clarification. *Social networks* 1(3), 215-239 (1978).
8. Mariani, J., Paroubek, P., Francopoulo, G., Hamon, O. : Rediscovering 15 years of discoveries in language resources and evaluation: The LREC anthology analysis. In: Proceedings of LREC, pp. 26-31 (2014).
9. Sheikahmadi, A., Nematbakhsh, M. A., Zareie, A.: Identification of influential users by neighbors in online social networks. *Physica A: Statistical Mechanics and its Applications* 486, 517-534 (2017).
10. Khadangi, E., Bagheri, A.: Presenting novel application-based centrality measures for finding important users based on their activities and social behavior. *Computers in Human Behavior* 73, 64-79 (2017).
11. Dave, K., Bhatt, R., Varma, V.: Identifying influencers in social networks. In: Proceedings of the 5th International Conference on Weblogs and Social Media, pp. 1-9 (2011).
12. Jabeur, L. B., Tamine, L., Boughanem, M.: Active microbloggers: Identifying influencers, leaders and discussers in microblogging networks. In: International Symposium on String Processing and Information Retrieval, pp. 111-117, Springer, Berlin (2012).
13. Gionis, A., Terzi, E., Tsaparas, P.: Opinion maximization in social networks. In: Proceedings of 2013 SIAM International Conference on Data Mining, 387-395 (2013).
14. Bigonha, C., Cardoso, T. N. C., Moro, M. M., Gonçalves, M. A., Almeida, V. A.: Sentiment-based influence detection on Twitter. *Journal of the Brazilian Computer Society* 18(3), 169-183 (2012).
15. Biran O., Rosenthal, S., Andreas, J., McKeown K., Rambow, O.: Detecting influencers in written online conversations. In: Proceedings of the Second Workshop on Language in Social Media, pp. 37-45, ACL (2012).
16. Quercia D., Ellis, J., Capra, L., Crowcroft, J. : In the mood for being influential on twitter. In: Proceedings of IEEE Third International Conference on Social Computing, pp. 307-314 (2011).
17. Hamzehei, A., Jiang, S., Koutra, D., Wong, R., Chen, F.: Topic-based Social Influence Measurement for Social Networks. *Australasian Journal of Information Systems* 21 (2017).
18. Weng, J., Lim E. P., Jiang, J., He, Q.: Twiterrank: Finding topic-sensitive influential twitterers. In: Proceedings of the Third ACM International Conference on Web Search and Data Mining, pp. 261-270, ACM (2010).

19. Katsimpras, G. Vogiatzis, D., Paliouras, G.: Determining influential users with supervised random walks. In: Proceedings of the 24th International Conference on World Wide Web, pp. 787-792, ACM, New York (2015).
20. Palau, R. M., Moens, M. F.: Argumentation mining: the detection, classification and structure of arguments in text. In: Proceedings of the 12th international conference on artificial intelligence and law (2009).
21. Amigó, E., Carrillo-de-Albornoz, J., Chugur, I., Corujo, A., Gonzalo, J., Meij, E., De Rijke, M., Spina, S.: Overview of RepLab 2014: Author Profiling and Reputation Dimensions for Online Reputation Management, In: International Conference of the Cross-Language Evaluation Forum for European Languages, pp. 307-322, Springer, Cham (2014).
22. Kleinberg, J. M.: Authoritative sources in a hyperlinked environment. *Journal of the ACM* 46(5), 604-632 (1999).
23. Lü, L., Zhang, Y. C., Yeung, C. H., & Zhou, T.: Leaders in social networks, the delicious case. *PloS one* 6(6) (2011).
24. Deturck, K.: Détection d'influenceurs dans des médias sociaux. In : Proceedings of TALN 2018, vol. 2, pp. 117-130, ATALA (2018).