Application of AI Techniques to Deep Web Social Network Analysis

Guillaume Gadek (1, 2), Stéphan Brunessaux (2) and Alexandre Pauchet (1)

(1) LITIS Normandie Université
76800 Saint-Étienne-du-Rouvray
FRANCE

(2) Airbus Defence and Space
78996 Elancourt Cedex
FRANCE

guillaume.gadek@airbus.com stephan.brunessaux@airbus.com alexandre.pauchet@litislab.fr

ABSTRACT

Online Social Networks (Twitter, Facebook, Telegram, and so on) have taken a huge place in the informational space, and are often used for propaganda, manipulation or even recruitment by authoritarian states or terrorist organisations.
As the quantity of information makes difficult the human exploitation, solutions to support the military decision makers can only come from the use of AI techniques to extract intelligence from posted texts, to qualify the users' behaviours, and to identify the emerging social structure.
In this article, we will illustrate how to use such techniques on a very peculiar social network hidden in the Dark Web. Thanks to the use of TOR, its members have an excessive feeling of freedom, enabling unconventional posts including hate speech and other hacker offers.
We propose to unfold an analysis of 1000 days of activities of this network using NLP techniques to find the most interesting emerging topics. On selected areas of interest, we explain how these techniques can also support the discovery of key actors of the network. We then proceed with a ML-based profiling of the user behaviors. Finally, we introduce influence and cohesion scores for groups of users, which help their characterization and evaluation.

1.0 INTRODUCTION

Web 2.0 and the social media have introduced a new paradigm in the informational space. Its specificities as a broadcast media enable anyone to be a new source of information, entertainment, or propaganda, at will. It gives power to the consumer, signalling defects on products; it gives the opportunity to the citizen to directly chat with his/her representatives; overall, the social web enables the user to be more connected and informed.

However, the presence of threats on Twitter and other social platforms is attested: the jihadist group Daesh is said to benefit from the support of 100,000 accounts on Twitter [1]; Russia is suspected by the FBI of interferences in the 2016 US election; in almost every election, worldwide, the intensive use of deception has been widely observed during the last years. Automation enables any party to produce and diffuse content quite effectively, if we consider the appearance of malicious accounts in the classic media a victory.

The diversity of threats implies a multiplicity of the tools to tackle all of the aspects of the problem. The quantity of messages exchanged in an Online Social Network can be astounding, with almost 500 million new tweets each day.

As is possible the automation of their publishing, it is also possible to automate the reading. Opinion mining and topic detection enable the analyst to grasp the trend on a mass of texts, providing a classification along various keywords and characteristics.

The user accounts can be classified according to various taxonomies. In this article we distinguish the profile (who the user claims to be), and the behaviour (how s/he performs actions). Unsupervised machine learning helps here to cluster the users along similar types, and enables the analyst to recognise similar situations.

The never ending list of “friendship” connections, as well as the list of interactions between user accounts, is fertile to grow graphs. This kind of graphs can be used to easily compute social influence scores [2], and to detect one of the most interesting features of a social network: people naturally gather in groups, where there is more continuous interaction. To detect these groups from graphs, we rely on community detection algorithms [3]; specific measures help the analyst to characterise the impact and the pertinence of such groups with regards to his preoccupations [4].

Galaxy2, on Tor, is a strange place. Created in 2015 to replace a predecessor social network (simply called Galaxy), it has been branded as the most frequented social network on Tor. In this article, we show the added value of a variety of tools to scan this very peculiar social network, analyse the topics of discussion during Galaxy2’s uptime, detect its key actors, and discover and characterise its communities. Finally, using all this information, we perform a case study, showing the interest of dedicated tools for social network analysis.

This paper is organised as follows: Section 2 presents the website and the context of TOR, Section 3 presents the tools used, Section 4 shows the case study; Section 5 proposes a discussion about our results and concludes this article.

2.0 PRESENTATION OF GALAXY2

2.1 The TOR Network

Partly developed through DARPA funding in the 90s, Tor, The Onion Router, has been launched in 2002. Aiming to bring anonymity to the data flow, it relies on an encryption and routing protocol, named onion, to hide the content of the packets to the transiting servers.

Tor is used either to anonymously access the clear web, as well as to access the so-called Dark Net: some websites are only accessible through the onion protocol, protecting the host. Well-known examples include a Wikileaks portal, and a disrupted illegal marketplace, the silk road.

As anyone can use the Tor network and enjoy the increased privacy, it helps to breach censorship and can be used by journalists, political activists or anyone else. Unfortunately, Tor is also used to host illegal content, such as hacker forums, drug marketplaces, dark forums, porn and pedo-pornography (https://www.fbi.gov/news/stories/playpen-creator-sentenced-to-30-years).

However, due to the perceived complexity of use and the strong incentive to link one’s Internet profile with his real life (shopping, all kinds of services...), the global usage of Tor is not very high. Moreover, some countries deploy firewalls, laws and rules to prohibit this kind of tools. Figure 1 shows the global usage of Tor as of 2014, sizing the countries by the total amount of users. Western Europe hosts most of them; Russia and the Middle East are also very present. On this map lies an invisible giant: continental China disapproves Tor.

2.2 Galaxy2: History and Main Features

Founded in 2015 after the disruption of a previous Tor hosted social network (“Galaxy”), Galaxy2 is based on an open source framework named elgg, enabling to build small social websites. According to its anonymous founder, Lameth, *“The server broke down and your terrible host here (me, not the current host, mind you!) hadn’t been keeping regular backups off the server”*. The service is shut down since end of October 2017 (https://socialmediaalternatives.org/archive/collections/show/10).

The main features of Galaxy2 include The Wire, a space for microblogging posts; blogs, polls and pages; image and file sharing. Because of the anonymity introduced by the Tor network, the users are less keen to use their real names, and do not expose personal data. Thus, only the private direct messages are kept private, and it was possible to consult all the posts, friendship connections, images and comments performed by the users over the almost three years of uptime.

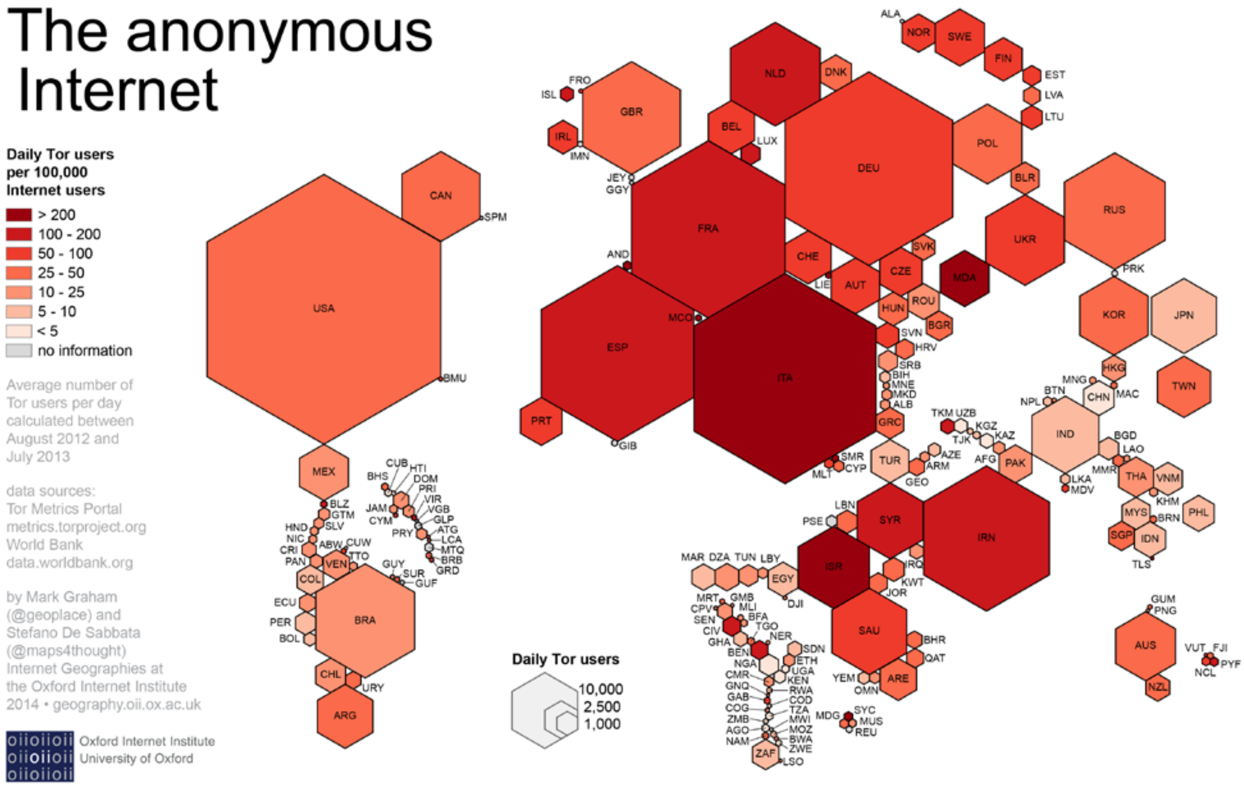

Figure 1: Geographies of Tor, Stefano De Sabbata.

The collection of our corpus covers a time range from Galaxy2 creation, on 9th of January, 2015 to the 22nd of September, 2017. It includes all traces of activity logged and then publicly available, but does not consists of personal profiles nor photos, excluded of the collect.

3.0 TOOLS FOR SOCIAL MEDIA ANALYSIS

3.1 Focused on Text Processing

The first axis of analysis deals with text processing. The tremendous amount of new messages emitted through the Online Social Networks pushed the development of machine text understanding and summary. Moreover, as users often express their feelings and opinions about products, facts and events, one of the most necessary tasks is opinion mining, which can briefly be explained as the automatic computation of a couple (topic, sentiment) matching a message [5].

The sentiment, or polarity, enables the evaluation of the support or reject of an entity by the author of a message. It can be either computed specifically with a topic, using a learning corpus, or generically using linguistic resources. Such a resource, SentiWordNet [6], associates polarities with sets of synonyms. An analyser, Vader [7], has been made available as an open source module, and compiles various resources and a classifier to compute sentiment scores, at document level.

Topic detection can be performed either through a supervised or an unsupervised method. An analyst can describe a topic using a few keywords, or rely on text clustering to let an algorithm cluster similar texts together, based on the semantics. Instead of matching pre-established categories, this approach adapts well to new corpora, enabling the analyst to grasp the topic distribution.

Often cited as a reference for text clustering, the Latent Semantic Analysis [8] can briefly be resumed as a dimensionality reduction applied on a matrix where each line consists of the tf*idf representation of a text of the training corpus. A similar method, LDA (Latent Dirichlet Allocation) [9] is said adapted to short text topic classification. It can be associated with a module of visualisation of the detected topics, to better depict their distribution [10].

3.2 User Profiling

One of the main tasks of social network analysis, or social listening, aims to spot the key actors, i.e. the most important users of the network. The focus is commonly directed either towards the influence [2, 11, 12], defined as the ability of a user to trigger actions from other users, or towards community detection [13, 14].

Influence is commonly defined as the ability to engage, i.e. to make the other users do something (typically, sending a message or clicking on a URL). A complete method has been exposed to query Twitter for follow links, to build a social graph and to update it [2]. This enables the computation of a daily PageRank, as a score of influence. However, using the public free API, the queries of follow links are slow, and a very long time (counted in months) is needed to obtain the complete social graph. To evaluate influence, various indicators, from the basic ones to the more complex, have been compared, among which the number of retweets, number of followers and PageRank computed on a graph of friendship links [11, 15].

Sometimes, the focus is not set on the most influential accounts, but on a general profiling to help the analyst understand “who” is a given account: what its typical behaviour is. Machine learning tools can answer this question using numeric features, through clustering or classification. When the data represent a graph, specific algorithms, such as RolX [16], enable a tagging of the nodes based on their graph features.

3.3 Group Detection

Studies in psychology highlight the power of the “group” as a structure, and the variety of its impacts on individual behaviour [17, 18, 19]. In computer science, a study on polarisation of the attitudes on online debates concludes similarly: groups can be considered as influencers, too [20].

3.3.1 Detection

To detect these groups, the main approach relies on community detection from friendship relation graphs. A community is a set of nodes, in a graph, more densely linked than a random graph. An intuitive way to discover groups is to look for communities in a graph built from the social network data.

However it has been pointed out that the social graph (“follow” graph) does not illustrate user interaction [21]. As a response, random walks community detection have been applied on a graph of interactions between Twitter users (retweets and mentions) [22].

To help choosing the best community detection method, a comparison of various state-of-the-art algorithms [23] on large networks (around 400k nodes) has been performed between FastGreedy [24], InfoMap [25] and Louvain/Blondel [3]. In this experiment, the Louvain method seems to be, by far, the fastest, for similar results in terms of quality.

3.3.2 Characterising the strength of the groups

Topological graph features such as the modularity value or the density of the detected groups can be considered to evaluate a community detection algorithm. A comprehensive review of scoring functions for community evaluation details the following measures [26].

Internal density (d) follows the intuition that communities are more dense than a random set of nodes. However, for large, non-complete communities, its value may be very low. On the other side, the smallest possible communities (pairs of nodes) have an internal density of 1.

Triad participation ratio (TPR) evaluates the number of nodes belonging to triads, or triangles, in community S. A value of 1 means that the group is strongly internally linked. Communities should include numerous triangles, showing interaction between the members.

Conductance (c), evaluates the quantity of edges linked with the other communities [27]. It illustrates the community behaviour, ranging from 0 (not linked towards the exterior) to 1 (strongly connected to other groups).

Modularity (Q) focuses on the number of internal edges (which link members of a community), compared to a “normal” quantity if the graph was random. Introduced by [28], a high modularity Q, close to 1, denotes a good partitioning at the graph level.

These topological metrics enable the analyst to check whether a community is sufficiently connected internally and/or with its boundaries: internal interaction and reach towards the exterior are measured.

3.3.3 Characterising the topical cohesion and relevance of the groups

To measure each of the obtained groups based on the topics expressed by the users, two topical metrics ξ and ρ, inspired from machine learning (ML) precision and recall, have been proposed [4]. These two ML inspired topical measures, which enable to better attribute a weight of the importance of a community on a given topic (ρ), and to evaluate the internal cohesion of a group (ξ), make the link between the graph and the semantic of the texts. Another measure (ϑf.igf) [4], similar to the natural language processing domain tf*idf, evaluates the specificity and the importance of a topic in a group. The analysis of these scores yields interesting insights on a social and textual corpus. For marketing purposes, one can identify then target a customer group, helping to induce a positive opinion about a new product.

4.0 CASE STUDY

In this section, the Galaxy2 network is exposed through the detailed explanation of the collected data type. Then, results are proposed to analyse the texts, the user accounts and the groups, which are an emerging social feature.

4.1 Types of actions logged

Table 1 presents the number of actions performed, per type, over the period of analysis. Actions are here regrouped to better grasp their diversity: first are the comments. Users can add some words to react to the other’s actions: publication of photos, files, new pages or polls, and of course blog posts.

Table 1: Description of the types of actions present in the corpus

	Type of action	Quantity
Comments	commented on	3
	commented on the album	149
	commented on the file	170
	commented on a page titled	192
	commented on a bookmark	239
	commented on the photo	375
	commented on the poll	672
	commented on the blog	4032
Files	uploaded the file	323
	created a new photo album	381
	added < some > photo(s)	516
	has a new avatar	1442
Creation	created a page	109
	created a poll	113
	added a new discussion topic	451
	created the group	548
	published a blog post	1944
Replies	voted on the poll	1328
	replied on the discussion topic	1340
Microblogging	posted on < a wall >	606
	posted to < The Wire >	29210
Connections	bookmarked	541
	joined the site	19233
	joined the group	26566
	is now a friend with	61027

A second kind of actions gathers the upload of files and images: those are deemed to be shared, and frequently include pictures from the Anonymous. As Tor and Galaxy2 philosophies are not based on personal holidays pictures publication, the quantity of images is quite low (around 600 pictures, in total). Avatars consist of profile picture change, though this feature has not been used a lot on Galaxy2.

Creation of pages, groups and polls and their modification (new topics of discussion) are seen as an aspect of the network life. This step of page creation is to be linked with the cluster of votes, replies and publication on said pages/groups.

Microblogging features were the predominant ones, with 29,000 posts on The Wire, the message feed. Message size was not limited, though frequently short.

Finally, the last cluster includes some connection information: account creation notifications ("joined the site"), link establishment between a user and a page, a group or another user (friendship).

4.2 Text: topic and sentiment repartition

4.2.1 LDA for topic detection

LDA is commonly used with a high number of topics; a number of 200-500 is usually recommended as a first try (http://radimrehurek.com/gensim/models/ldamodel.html). As a trade-off between computation time and representativeness of the topics, a number of 40 topics is chosen. Compared to other values, this choice results in a distribution of the documents over the topics, also limiting the number of too-close topics.

Figure 2 shows a tool, proposed by [10], which enables to navigate along the 40 detected topics through the word frequencies in a given topic (in red), compared to word frequencies over the whole corpus (in blue), in the right panel.

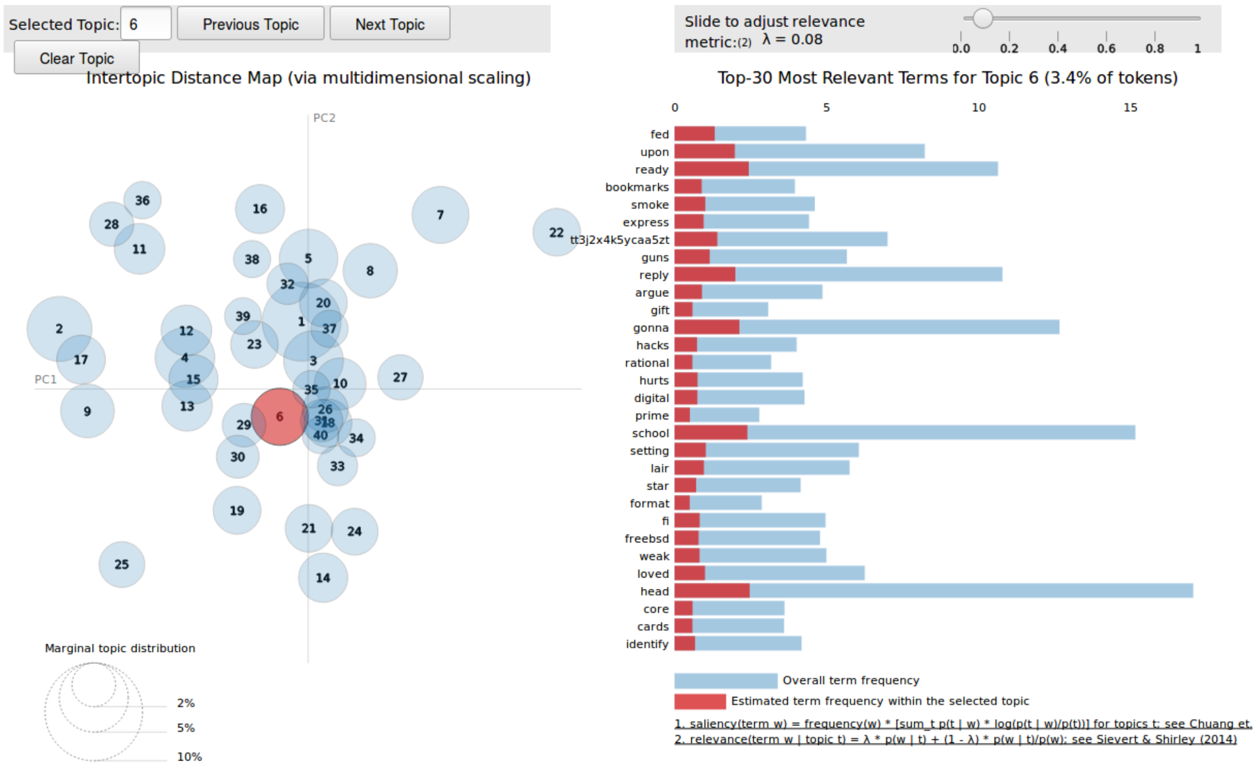

Figure 2: Visualisation of the topics

A principal component analysis (PCA) is performed to place the topics in a 2D-space, in the left of the diagram. The area of the circles is proportional to the quantity of documents clustered in the topics. Thanks to this tool, the analyst can visualise the topic distribution of, here, more than 30,000 documents, while having the ability to see which words are more frequent or more relevant in each of the topics.

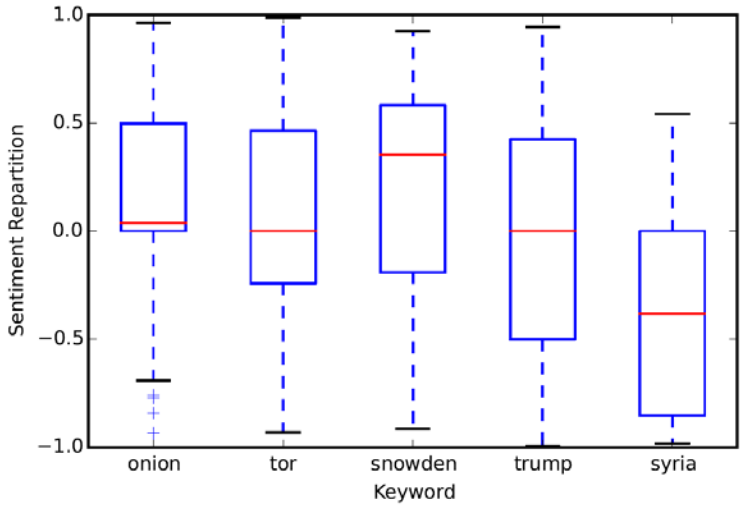

Figure 3: Sentiment distribution for some keywords

4.2.2 Sentiment exploitation

Each text is also analysed with regard to the sentiment, representing the overall polarity of a text, as perceived by a reader. Even if some of the sentiment predictions may be inaccurate, it results in a global view. For instance, Figure 3 proposes to compare the polarity of the texts containing some keywords: onion, Tor, Snowden, Trump and Syria (no distinction has been made between uppercase and lowercase). Instead of an average value, and as one message is enough to set either a max or min sentiment value, we represent it as a boxplot, showing the statistical distribution. Onion is globally well perceived, with positive messages. Edward Snowden seems to enjoy a good reputation, while Galaxy2 members are more mitigated about Donald Trump. Finally, the polarity is more negative about Syria, because these messages are susceptible to mention the current conflict.

The sentiment view can also be exploited to characterise a user activity. As an example, Figure 4 exposes the sentiment repartition of the messages of two active authors: XL33t and Fenris. On both figures, an artefact, for a sentiment of 0.0, exists, representing the messages where no polarity could be extracted by the analyser (no known sentimental word). However, the repartition shows some signals. On the first figure, XL33t seems to mostly emit very positive messages. The second author, Fenris, covers the whole range of sentiment, avoiding to publish too positive messages. This view let the analyst grasp one aspect of a user behaviour, as well as to estimate his activity (thanks to the number of messages, indicated on the y-axis).

4.3 Influential users and types of behaviour

4.3.1 Top-5 key users

Influence can be measured through a variety of prisms. In a first glance on a social network, one can look for the key actors, detecting a “top5”. Table 2 compares the most connected users (having the highest number of friends, noted #Friends); the most popular users (the most mentioned, noted #Mentions); the users of reference (a score based on the mentions graph, in the column Sc(Mentions), explained in the following paragraph); and the most active ones (performing the highest number of actions on Galaxy2).

Inspired from the literature [11, 15], a score of influence is computed from Gm the graph of mentions. The intuition is the following: a mention gives some value, some social capital to the mentioned user; this value is higher if the author of the mention already has some influence him/herself. The PageRank algorithm also follows this intuition and has been shown to empirically match this representation of influence. In consequence, ScoreMentions is the ranking resulting from the Pagerank scores of each user-node in Gm.

Table 2: Top5 of the Users Along Three Dimensions

#Friends		#Mentions		Sc(Mentions)	#Actions	
XsyntaX	8187	XsyntaX	514	Oxyy	XsyntaX	7974
prozac	3182	Oxyy	164	XsyntaX	kheper	1846
Spooky	3012	Lameth	121	Lameth	prozac	1783
xl33t	2140	Fenris	119	ChatTor	ChatTor	1583
kheper	1922	cpnnemo	96	MahaKali	Spooky	1569

In a few words about Table 2: XsyntaX outperforms every other account, with 7974 actions performed during the scope of the study. The second, kheper, performed 1846 actions. Lameth is the founder of this social network; as such he is often mentioned, either for thankful messages for hosting, or for administration purposes. ChatTor provides a chat service through Tor; this account promotes the service and publishes news and updates.

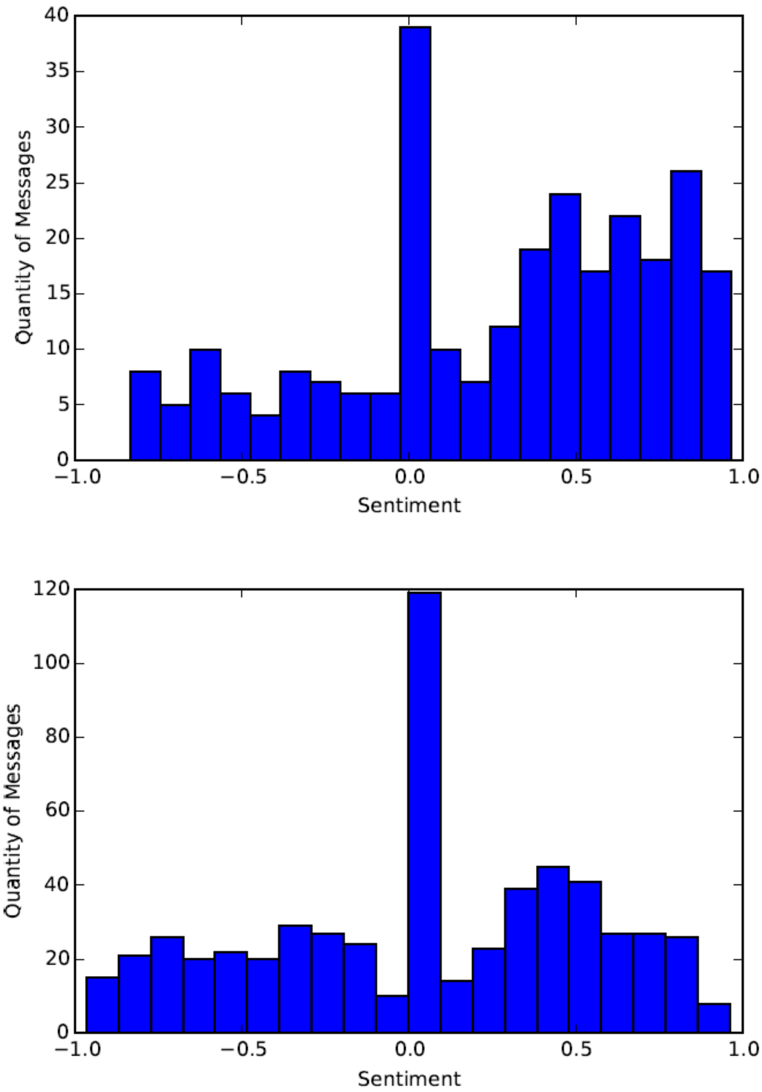

Figure 4: Visualisation of the sentiment emitted by two users: XL33t and Fenris

Although some names are present on every influence score Top5, they do not all indicate the same kind of influence. Lameth is mentioned a lot, due to his status of host. However, he did not intend to be the friend of everyone else, or to be the most active. Xl33t connected to many accounts, but is not mentioned as much as the other influential users. Influence and key-actors take many different forms, while only the symptoms can be measured.

4.3.2 Types of behaviour

The social network analyst cannot be expected to know who every user on a platform is. More precisely, the data of interest can be narrowed along a few aspects to grasp the essence of the behaviour.

Biography contains the user identity: username, ID, date of creation and other available data. Style explains how the user writes: length of messages and quantity of punctuation usually make a difference, and can be combined with the main topic of interest and global sentiment polarity. The social aspect details the number of friends, of mentions, and the resulting influence of the user, both from the graph of friendship and the graph of mentions. Media groups the features explaining the type of actions and objects posted (be it photos, texts, etc). Finally, the temporal aspect covers the global rhythm of publication of the user: average number of posts on TheWire per day, and average number of actions per day.

All these numerical features enable the repartition of the users along a few typical clusters, whose reduced number (commonly 4 to 6, depending of the data) enables the analyst to get used to these unsupervised labels.

Technically, this repartition is made in a few steps. First the data is cleaned, so as to avoid zero divisions and exponential distribution of the data: some features are converted to their log-values to reduce dispersion, and normalised to follow a 0-mean, 1-standard deviation distribution. A PCA reduces the dimension of the problem, switching from 34 variables to 5 and still keeping most of the dataset variance. Finally, a k-means finds the groups; a number of 4 groups is a good trade-off between dataset dispersion and cluster shape.

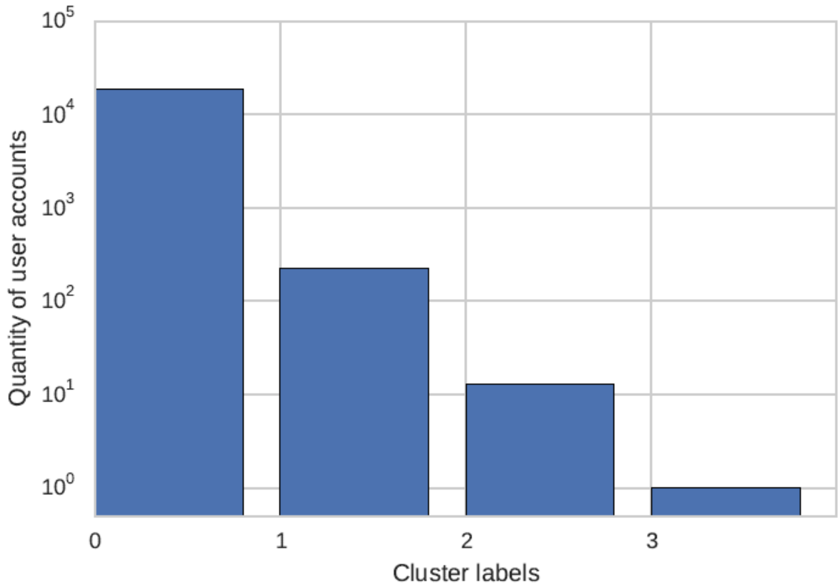

Figure 5: Repartition of the user profiles

Figure 5 shows the repartition of the users along the 4 types of profiles. The first type (0) is the most common, including more than 90% of the profiles: most of the accounts are created, perform one or two actions, and then are forgotten. A second type, profile (1), characterises the regularly active users, even though they are not the central users of the network. Finally, profile (2) and the sole user in profile (3) are the central users, very active, those who produce most of the contents on the network. A similar repartition had already been remarked on networks such as Twitter [11].

To picture the shape of the profile clusters, Figure 6 represents each of the 19,177 Galaxy2 users along the first two dimensions of the PCA. These dimensions are linear combinations of the behavioural features, and consequently are quite abstract. However, positions near (0,0) are linked with very low levels of activities (the red cluster (0) being the most populated). In blue, cluster (1) shows the many different ways to be a standard active user. The green cluster (2) is very dispersed but matches high levels of publications, be it by posts, comments or creation of pages. Finally, the purple cluster (3) is constituted of only one account, Spooky, performing an average of 2.5 actions each day.

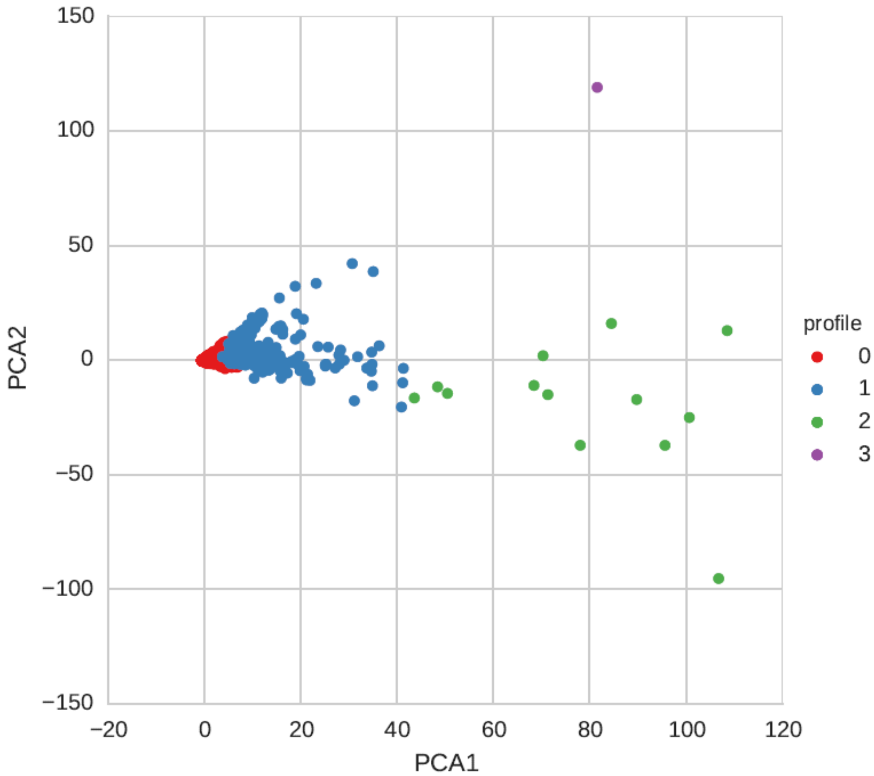

Figure 6: Visualisation the clusters of users in a PCA space

When the analyst reviews the posts, thanks to this step of profiling, he can quickly evaluate if the messages comes from well-established, influential users, or whether a given topic is dominated by usually inactive accounts, or show a unusual repartition.

4.4 Interactive groups, their cohesion and structure

4.4.1 About the user mentions

A common practice on social networks consists to mention people, using an “@” sign before their usernames in the text of the messages. At first this was only a social behaviour without any technical feature in the code of the social platform; nowadays most platforms recognise such mentions.

As Galaxy2 is based on the old-fashion social framework elgg, such a feature is not hard-coded, but is still used by some members. Following the intuition that this interaction (the mention) would bring us valuable information, a mention extraction module has parsed each of the publications in search for “@” signs followed by usernames.

From the set of mentions M associated with the authors A, a graph of mentions, Gm = (V,E) is built. V is the set of the union of the authors and the mentioned users; edges e ϵ E links the author of a message to the eventual mentioned user. This graph Gm is composed of 968 nodes, linked by 2342 edges (there was a total of 5481 mentions: some accounts have been mentioned by a same user various times). The Louvain community detection [3] results on a set of 31 communities, amongst which only 11 contain more than 3 members.

Figure 7 illustrates one of these communities. Red nodes are part of the detected community; blue nodes are the only external contacts, the boundary of the community in the graph GM. A central user, Bishop, is the target of a few mentions, notably from an audacious account, Nishikino Maki (he claimed to be “building a porn site cuz i can”). In a glance, we see that this community is centred on Bishop, who is not the most active in terms of community life (he does not mention the other members of the group).

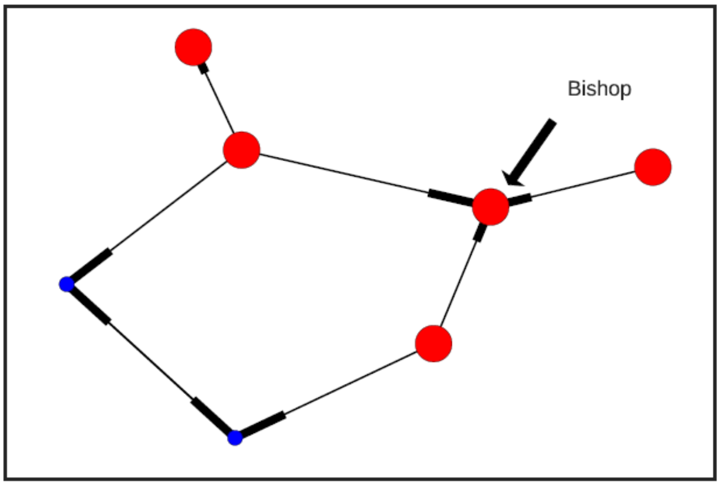

Figure 7: Visualisation of a small community and its contacts through mentions

This community view helps the analyst to see the relations between a user and the rest of the network, and can guide his exploration along the interaction between accounts. However, as this graph of mentions is built from a social usage (using “@” to mention a user), the tool probably missed some differently written mentions and thus may be incomplete (this graph Gm is one of the graphs of interaction, not the only one).

4.4.2 The friendship graph

The most known feature of social networks is the friendship connection. On Galaxy2, one can claim his friendship with another user, which is considered reciprocal on Galaxy2. This action is part of the list of the collected activities, and enables to link users along these friendship relations. Let Gf be the graph of friendship, where nodes are user accounts. Edges represent a claim of friendship.

From this graph, communities emerge, using the Louvain algorithm to reveal them. These communities may vary in number, based on the algorithm used. Here, 32 communities were obtained, amongst which 11 include more than 3 users. Figure 8 shows the distribution of the communities along two topological measures: x-axis gives the size of the group while y-axis represents the conductance, which is the proportion of boundary edges, linking the community to its environment, showing either its influence or its isolation.

On Figure 8, a group attracts the attention: with size=125 and conductance=0.74, this community is less tied to the rest of the network than it normally should. Its topical scores, based on [4], require some interpretation. ξ = 0.064 means that only 6.4% of its members, as a maximum, have been active about the same topic; however, compared with the duration covered on the corpus, this situation is not surprising. In particular, this "main topic" clusters mainly around two sets of texts at different dates. The first refers to a small pornographic discussion group; the second concerns the use of peer-to-peer mail through Tor.

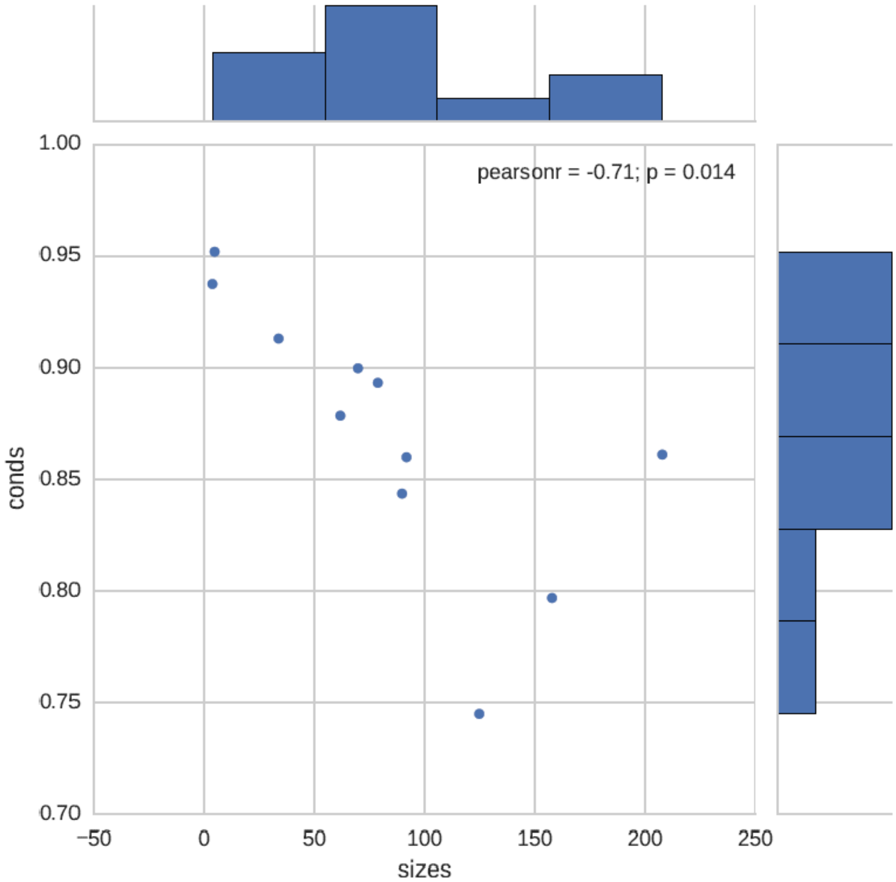

Figure 8: Visualisation of the topological features of the communities

These topical measures may benefit from improvements on the topic detection part, notably covering the temporality of topics and their labelling. This would lead to a more usable topic modelling and a more precise group characterisation for the analyst.

5.0 CONCLUSION

The emergence of modern social networks reveals the inherent complexity surrounding us. Although most of the analysis is focused on the most known Online Social Networks (OSN), such as Twitter or Facebook, other instances exist; Galaxy2 is much smaller than Twitter, but big enough to turn its analysis tedious due to the amount of exchanged messages.

We presented a whole system capable of analysing any type of OSN, showing its capabilities on a Tor-hosted social network. The functionalities include text analysis for topic and sentiment, user profiling and graph analysis, both on the friendship links as well as on the interaction between users. We showed that the tools, mostly developed for Twitter, can also be exploited for smaller networks. This approach has easily been adapted to the specific links and interactions of Galaxy2, as it could be for other platforms such as Reddit or Facebook.

Future work will be oriented on real-time processing, to analyse a social network as it lives. We have the intuition that a stream of messages would better fit the requirements in scaling and response times. We also desire to publicly release the Galaxy2 corpus, as one of the first complete social network archives.

REFERENCES

[1] G. Bertolin, Digital Hydra: Security Implications of False Information Online. NATO Strategic Communications Centre of Excellence, 2017.

[2] P. Noordhuis, M. Heijkoop, and A. Lazovik, "Mining twitter in the cloud: A case study," in Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on. IEEE, 2010, pp. 107–114.

[3] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," Journal of statistical mechanics: theory and experiment, vol. 2008, no. 10, p. P10008, 2008.

[4] G. Gadek, A. Pauchet, N. Malandain, K. Khelif, L. Vercouter, and S. Brunessaux, "Measures for topical cohesion of user communities on twitter," in Proceedings of the International Conference on Web Intelligence. ACM, 2017, pp. 211–218.

[5] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up?: sentiment classification using machine learning techniques," in Proceedings of the ACL- 02 conference on Empirical methods in natural language processing- Volume 10. Association for Computational Linguistics, 2002, pp. 79– 86.

[6] S. Baccianella, A. Esuli, and F. Sebastiani, "Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining." in LREC, vol. 10, 2010, pp. 2200–2204.

[7] C. J. Hutto and E. Gilbert, "Vader: A parsimonious rule-based model for sentiment analysis of social media text," in Eighth International AAAI Conference on Weblogs and Social Media, 2014.

[8] T. K. Landauer, P. W. Foltz, and D. Laham, "An introduction to latent semantic analysis," Discourse processes, vol. 25, no. 2-3, pp. 259–284, 1998.

[9] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," Journal of machine Learning research, vol. 3, no. Jan, pp. 993–1022, 2003.

[10] C. Sievert and K. E. Shirley, "Ldavis: A method for visualizing and interpreting topics," in Proceedings of the workshop on interactive language learning, visualization, and interfaces, 2014, pp. 63–70.

[11] H. Kwak, C. Lee, H. Park, and S. Moon, "What is twitter, a social network or a news media?" in Proceedings of the 19th international conference on World wide web. ACM, 2010, pp. 591–600.

[12] A. Rao, N. Spasojevic, Z. Li, and T. DSouza, "Klout score: Measuring influence across multiple social networks," in Big Data (Big Data), 2015 IEEE International Conference on. IEEE, 2015, pp. 2282–2289.

[13] A. Java, X. Song, T. Finin, and B. Tseng, "Why we twitter: understanding microblogging usage and communities," in Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis. ACM, 2007, pp. 56–65.

[14] T. Aynaud and J.-L. Guillaume, "Static community detection algorithms for evolving networks," in Modeling and optimization in mobile, ad hoc and wireless networks (WiOpt), 2010 proceedings of the 8th international symposium on. IEEE, 2010, pp. 513–519.

[15] C. Lee, H. Kwak, H. Park, and S. Moon, "Finding influentials based on the temporal order of information adoption in twitter," in Proceedings of the 19th international conference on World wide

web. ACM, 2010, pp. 1137–1138.

[16] K. Henderson, B. Gallagher, T. Eliassi-Rad, H. Tong, S. Basu, L. Akoglu, D. Koutra, C. Faloutsos, and L. Li, “Rolx: structural role extraction & mining in large graphs,” in Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, 2012, pp. 1231–1239.

[17] S. Milgram, “Group pressure and action against a person.” The Journal of Abnormal and Social Psychology, vol. 69, no. 2, p. 137, 1964.

[18] S. Milgram., “Liberating effects of group pressure.” Journal of Personality and Social Psychology, vol. 1, no. 2, p. 127, 1965.

[19] G. L. Cohen, “Party over policy: The dominating impact of group influence on political beliefs.” Journal of personality and social psychology, vol. 85, no. 5, p. 808, 2003.

[20] S. Yardi and D. Boyd, “Dynamic debates: An analysis of group polarization over time on twitter,” Bulletin of Science, Technology & Society, vol. 30, no. 5, pp. 316–327, 2010.

[21] K. H. Lim and A. Datta, “An interaction-based approach to detecting highly interactive twitter communities using tweeting links,” in Web Intelligence, vol. 14, no. 1. IOS Press, 2016, pp. 1–15.

[22] Y. Tyshchuk, W. A. Wallace, H. Li, H. Ji, and S. E. Kase, “The nature of communications and emerging communities on twitter following the 2013 syria sarin gas attacks,” in Intelligence and Security Informatics Conference (JISIC), 2014 IEEE Joint. IEEE, 2014, pp. 41–47.

[23] R. Cazabet and F. Amblard, “Simulate to detect: a multi-agent system for community detection,” in Web Intelligence and Intelligent Agent Technology (WI-IAT), 2011 IEEE/WIC/ACM International Conference on, vol. 2. IEEE, 2011, pp. 402–408.

[24] A. Clauset, M. E. Newman, and C. Moore, “Finding community structure in very large networks,” Physical review E, vol. 70, no. 6, p. 066111, 2004.

[25] M. Rosvall and C. Bergstrom, “Maps of information flow reveal community structure in complex networks,” Technical report, Tech. Rep., 2007.

[26] J. Yang and J. Leskovec, “Defining and evaluating network communities based on ground-truth,” Knowledge and Information Systems, vol. 42, no. 1, pp. 181–213, 2015.

[27] J. Shi and J. Malik, “Normalized cuts and image segmentation,” IEEE Transactions on pattern analysis and machine intelligence, vol. 22, no. 8, pp. 888–905, 2000.

[28] M. E. Newman, “Modularity and community structure in networks,” Proceedings of the national academy of sciences, vol. 103, no. 23, pp. 8577–8582, 2006.