Funded by the Horizon 2020 Framework
Program of the European Union
SocialTruth - Grant Agreement 825477

Deliverable D3.1

Social Mining Descriptors

Dissemination Level:	**PU**
Nature of the Deliverable:	**R**
Date:	**20/07/2020**
Distribution:	**WP3**
Editors:	**QWANT**
Reviewers:	**LSBU, UTP**
Contributors:	**ALL**

Abstract:

The main goal of this deliverable is to present and describe the social mining descriptors. It presents what will be made available at this point to the SocialTruth platform from Qwant in term of news and tweets.

***Dissemination Level:**	*PU= Public, RE= Restricted to a group specified by the Consortium, PP= Restricted to other program participants (including the Commission services), CO= Confidential, only for members of the Consortium (including the Commission services)*
**** Nature of the Deliverable:**	*P= Prototype, R= Report, S= Specification, T= Tool, O= Other*

Disclaimer

This document contains material which is copyright of certain SocialTruth consortium parties. All SocialTruth consortium parties have agreed to the full publication of this document.

Neither the SocialTruth consortium as a whole, nor any certain party of the SocialTruth consortium warrants that the information contained in this document is capable of use, or that use of the information is free from risk, and accepts no liability for loss or damage suffered by any person using the information.

The contents of this document are the sole responsibility of the SocialTruth consortium and can in no way be taken to reflect the views of the European Commission. The European Commission is not responsible for any use that may be made of the information it contains.

The commercial use of any information contained in this document requires a license from the proprietor of that information. For information and permission requests, contact the SocialTruth project coordinator Dr. Konstantinos Demestichas (ICCS) at cdemest@cn.ntua.gr.

The content of this document may be freely distributed, reproduced or copied as content in the public domain, for non-commercial purposes, at the following conditions:

a) it is requested that in any subsequent use of this work the SocialTruth project is given appropriate acknowledgement with the following suggested citation:

“Deliverable 3.1 Social Mining Descriptors (2019)” produced under the SocialTruth project, which has received funding from the European Union’s Horizon2020 Programme for research and innovation under grant agreement No.724087. Available at: http://www.socialtruth.eu“

this document may contain material, information, text, and/or images created and/or prepared by individuals or institutions external to the Socialtruth consortium, that may be protected by copyright. These sources are mentioned in the “References” section, in captions and in footnotes. Users must seek permission from the copyright owner(s) to use this material.

Revision History

Date	Rev.	Description	Partner
05/07/2019	0.1	First version of DDP	QWANT
12/07/2019	0.3	Initial deliverable structure (TOC)	QWANT, ICCS
29/07/2019	0.4	Initial deliverable content	QWANT
2/09/2019	0.7	Aggregation of partners	QWANT
18/09/2019	0.8	Updated version after periodic technical meeting (remotely)	QWANT
24/09/2019	0.9	Updated version after periodic technical meeting (remotely)	QWANT, ESF
28/09/2019	1.0	Updated version, based on peer-review feedback	QWANT
29/09/2019	1.1	Final version for submission	QWANT, ICCS
25/05/2020	1.2	Preparation of improvements, based on reviewers’ comments	QWANT
20/07/2020	2.1	Second version ready	QWANT

List of Authors

Partner	Author
QWANT	Stan ASSIER, Christophe SERVAN, Odilon GHEDIN, Pierre VIGNAUX
ESF	Sonia Collada, Nahid Oulmi
ICCS	George Koutalieris, Kostas Demestichas

Table of Contents

Index of figures

Index of tables

Glossary

WP	Work Package
NN	Neural Networks
GDPR	General Data Protection Regulation
API	Application Programming Interface

Executive Summary

This document describes the initial design of SocialTruth's social mining system and the social mining descriptors. The social mining system will be provided by Qwant as a service fully open to the SocialTruth platform. It will provide news and tweets described with meta data specific to the project. The social mining system is not meant to detect fake tweets by itself but to provide the SocialTruth platform with news and tweets which will include descriptors that will endow the SocialTruth platform with the detection of fake news features. The social mining system, within the SocialTruth project, is a provider of raw material collected for a specific purpose and enhanced in order to allow the SocialTruth platform to perform fake news detection.

The social mining system will be used for two purposes: to provide news and tweets upon request of the SocialTruth platform on a day to day operation basis; and, to gather news and tweets in order to form datasets which will be used to train the SocialTruth Neural Network models used in tasks 3.2 and 3.3. For the former objective the standard tweet descriptors coming from Tweeter will be enriched with a fine-grained sentiment descriptor extracted by Qwant from the tweet content analysis and with a retweet pace descriptor that will be computed from tweet's meta data. For the later objective, datasets will be provided only once. 684,473 fake political and gossips posts and 1,276,204 factchecked political and gossips posts have already been provided, these two datasets will be grow to up to 2 M tweets each by end 2020. Regarding sentiment analysis recent work shown its interest in fake tweet detection. Hence, we will add fined grained sentiment analysis (strongly negative, weakly negative, neutral, weakly positive, strongly positive) to allow an even better leverage of sentiment analysis for fake news detection.

1. Introduction

To proceed to fake-news detection the SocialTruth platform will use, among others, news from the Web and Tweets enriched with meta data. Qwant will be the one providing these inputs to the SocialTruth platform. This deliverable is focused on the description of how this will be achieved.

1.1 Forewords

Based and designed in Europe, Qwant is the first search engine which protects its users' freedoms and ensures that the digital ecosystem remains healthy. Our keywords are: privacy and neutrality. As such Qwant USED TO embeds a window on the social Web by integrating results from Twitter, as shown in Figure 1.

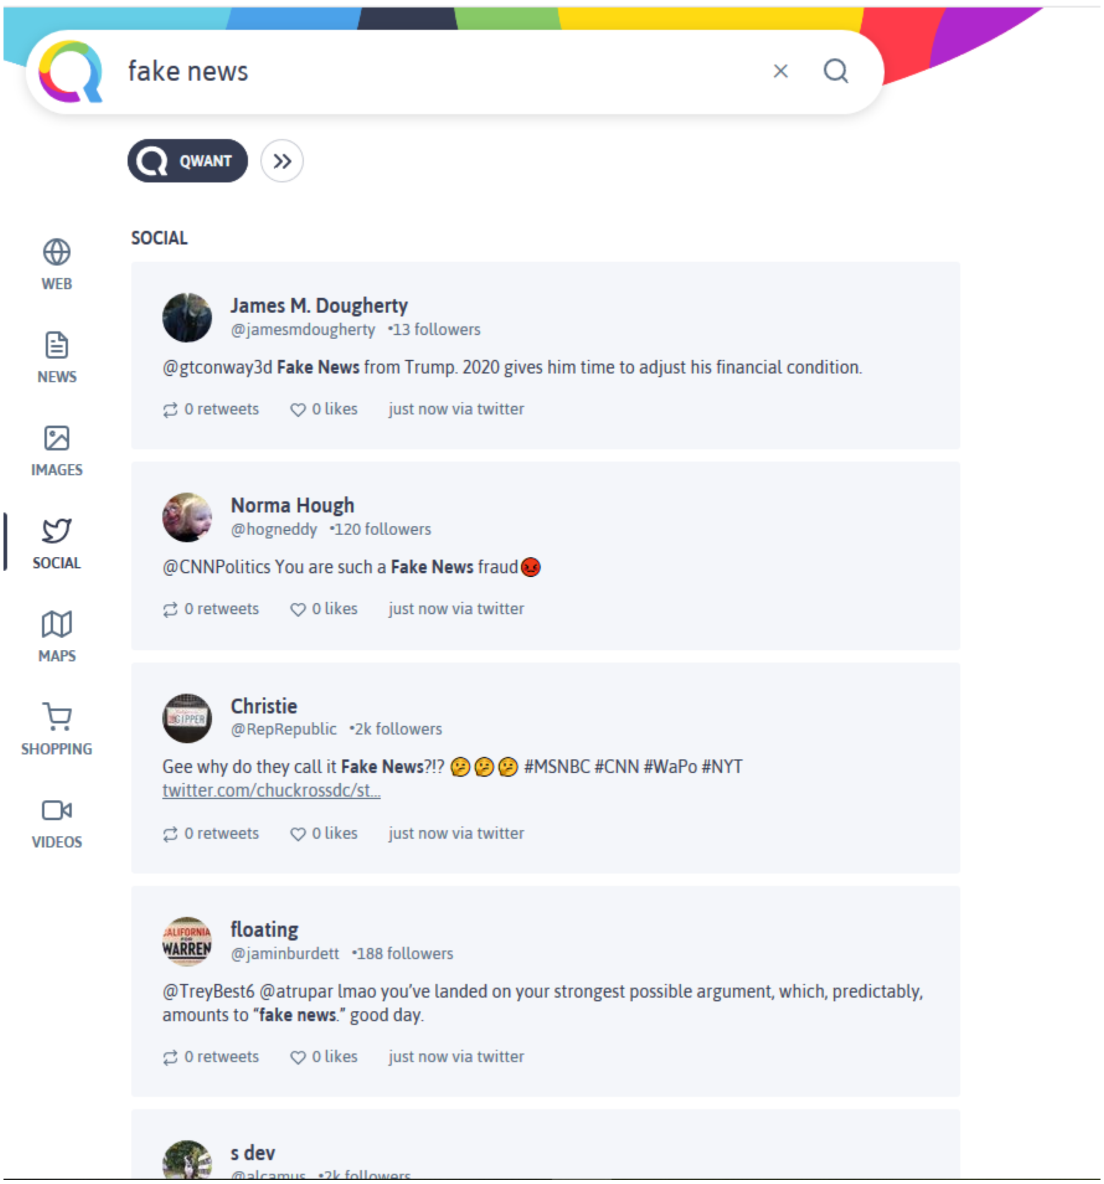

Figure 1: Qwant Screenshot with social results

Due to its strategic reorganization, Qwant does not include any more Tweets in its search engine results. Although the discussions between Qwant and Twitter have not concluded on this matter, currently, for this reason, a solution to cover explicitly the needs of the Socialtruth project is foreseen, as described in Section 3.1. Any further progress on the matter will be duly reported to the SocialTruth coordination team and the EC.

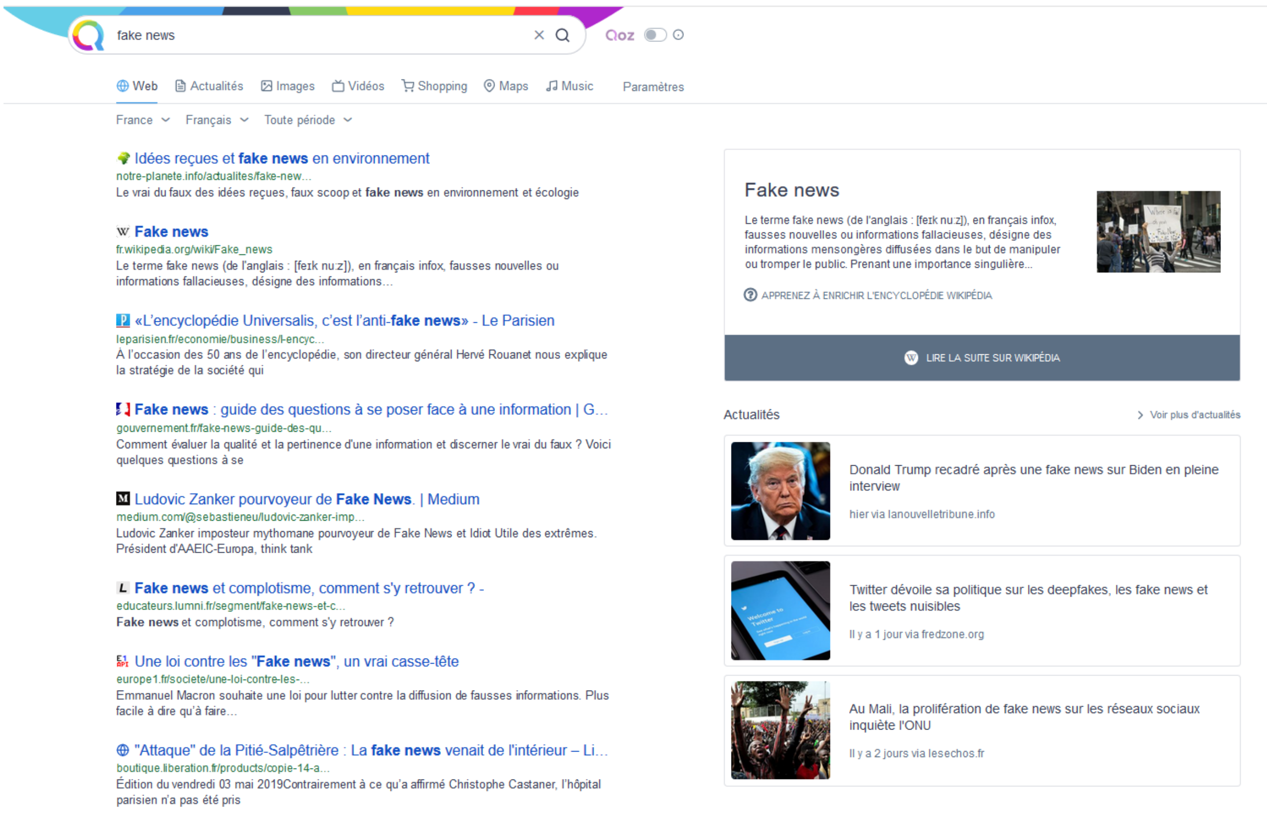

Figure 2: Qwant Screenshot without social results

1.2 Contribution of this deliverable to the SocialTruth solution

Qwant provides the required tools and expertise to collect news and tweets that will contribute the social mining aspects of the verification services that SocialTruth will deliver. Qwant will not analyze news nor tweet contents for fakeness detection but only collect and enrich the latter with fine grained sentiment and retweet pace upon the SocialTruth platform request and send them back to the platform. Also, Qwant will collect news and tweets to form the required datasets that will be used by the project members to train their models. A summary of the SocialTruth development activities and the associations between the components foreseen in all the project development WPs is presented in Figure 3 hereunder.

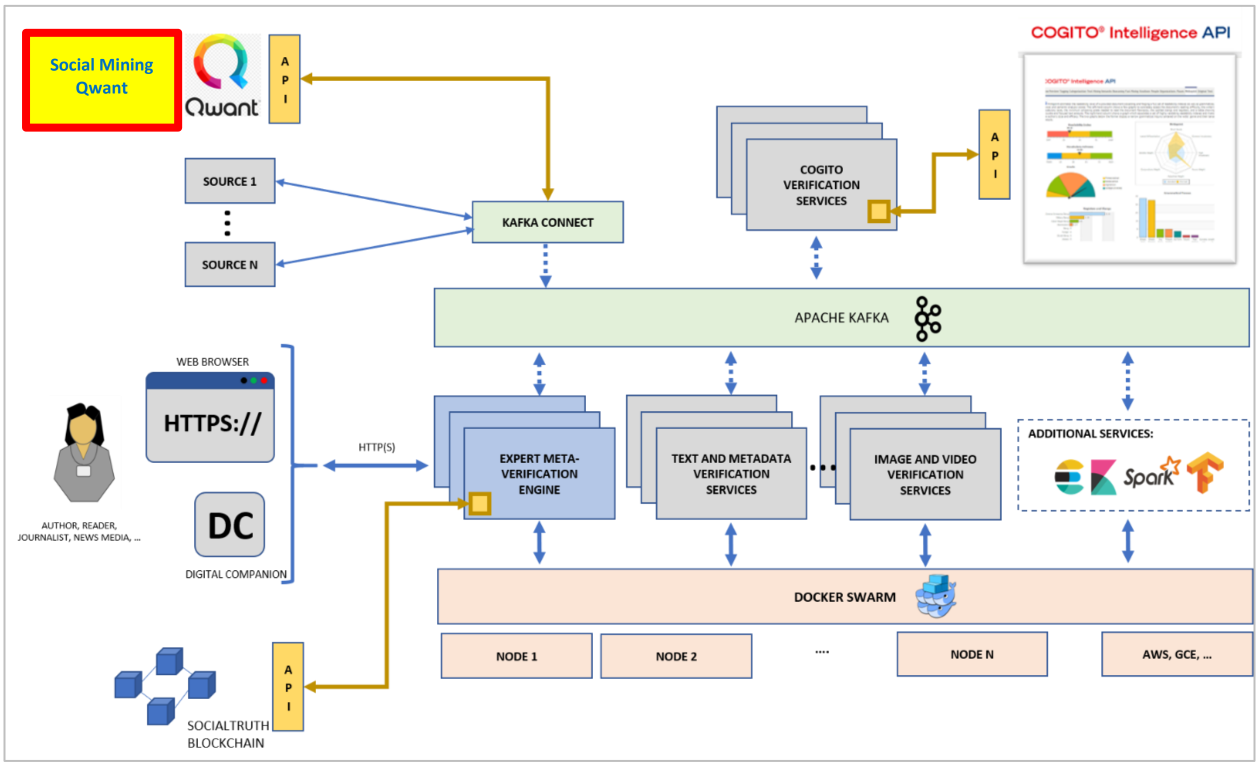

Figure 3: The Social Mining System within SocialTruth system

2. System description

2.1 General overview

The system is based upon APIs which allow the SocialTruth platform to request news and tweets from Qwant. The Qwant processes for setting the request performing the queries and provide the results, use internal Qwant resources for security, communication, computing and temporary storage. Also, Qwant processes, ensure full compliance with GDPR

2.2 News collection

For its search engine Qwant has developed a service which collect news. This service will be made available for the SocialTruth platform within the project. The communication between the SocialTruth platform and Qwant will use RESTful API. Here are the descriptions of the fields for the request of News.

Name	Type	Description	Example	Default value
q *	String (query)	Query expression	Fakenews	NA
Locale	String(query)	Search region locale	Fr_FR	fr_FR
Count	Integer (query)	Number of results to return		10
Offset	Integer (query)	Number of results to skip		0
safesearch	Integer (query)	Filter for results safeness (adult content, violence, porn) Available values: 0, 1, 2		&
Freshness	String(query)	Filter for results freshness	All / hour / day / week / month	all
Order	String(query)	Orders results	Relevance / date	relevance
Source	String(query)	Filter results by source	www.lemonde.fr	Relevance

Table 1: Qwant API fields description for News collection

The News will be sent back to the SocialTruth platform. They will be returned in JSON format via the Qwant API. Format is :

Name	Type	Description	Example	Default value
Status	String	Status	Success	NA
Data	String	News	NA	NA

Table 2: News description

2.3 Tweets collection and enrichment

The system is based upon two APIs invoking and two discreet Qwant services (Figure 4). The APIs are used in order for Qwant to communicate with the SocialTruth platform through the Qwant API and to communicate with the Twitter through the Twitter API. The Qwant processes for setting the request to Twitter and for performing the fine grained sentiment analysis and retweet pace computing use internal Qwant resources for security, communication, computing and temporary storage. Also, Qwant processes, ensure full compliance with GDPR.

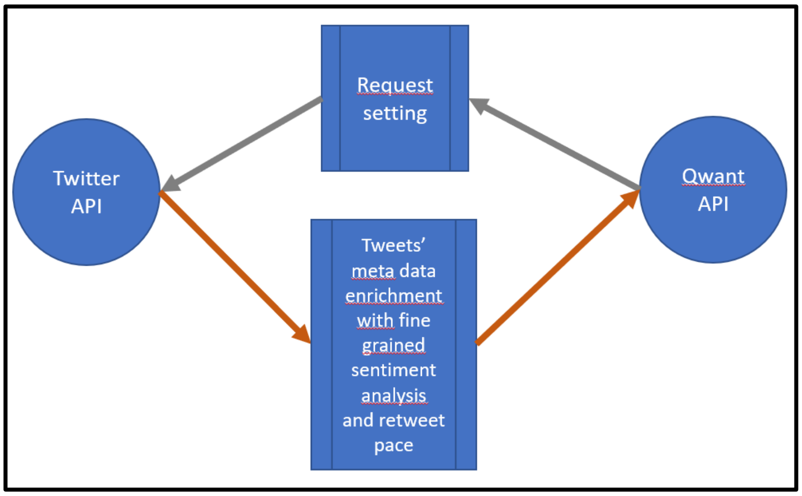

Figure 4: Tweets collection and enrichment

2.4 Operation descriptions

The SocialTruth platform sends its request and gets the tweets through the Qwant API and then the Qwant system accesses the tweets through its Twitter API. Before sending back the tweets to the SocialTruth platform Qwant performs a fine grained sentiment analysis and meta data calculation to enrich them. The following operation are carried out upon request from the SocialTruth platform (Figure 5).

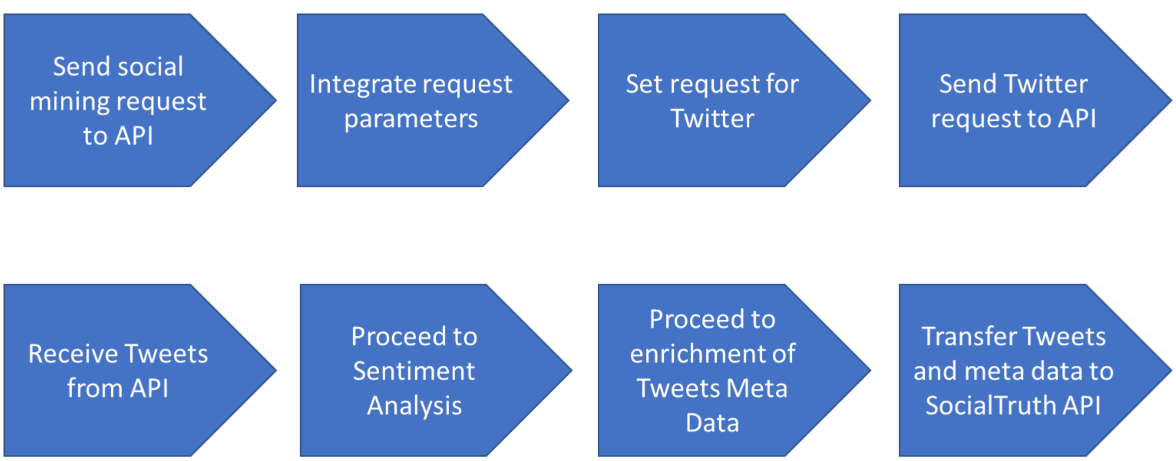

Figure 5: Operations description

The communication between the SocialTruth platform and the Social Mining system will use RESTful API. Here (Table 3) are the descriptions of the fields for the request of Tweets.

Name	Type	Description	Example	Default value
q *	String (query)	Query expression	Fakenews	NA
locale	String(query)	Search region locale	Fr_FR	fr_FR
Count	Integer (query)	Number of results to return		10
offset	Integer (query)	Number of results to skip		0
safesearch	Integer (query)	Filter for results safeness (adult content, violence, porn) Available values: 0, 1, 2		&

Table 3: Qwant API fields description for Tweets collection

The tweets will be sent back to the SocialTruth platform. They will be returned in JSON format via the Qwant API. Tweet content is described hereafter as Social Mining Descriptors.

2.5 Tweets Descriptors

The project will collect posts (tweets) from Twitter. The descriptors (Table 4) which will describe the tweets are:

Tweet descriptors	
Tag	**Description**
tweet_id	ID of the tweet
tweet_url	URL of the tweet's Web page
account_name	Username of the Tweetos
description	Content of the tweet
retweets	Number of times the tweet has been retweeted
likes	Number of times the tweet has been liked
lang	language of the tweet
followers	Number of followers of the Tweetos
Sentiment	Very Positive, Positive, Neutral, Negative or Very negative content of the Tweet
Retweet pace	Pace on which the considered tweet has been retweeted

Table 4: Tweet description

Within these descriptors, the "sentiment" one is a specific descriptor for the SocialTruth project. As the project progressed and the Consortium provided feedback, Qwant adds the retweet pace descriptors that could add to the accuracy of fakeness detection by the SocialTruth platform.

2.6 Retweet pace calculation

Retweet pace is known to be a very accurate indicator of fakeness regarding tweets [Zubiaga and Ji, 2014, Atodiresei et al, 2018]. Hence, as this information is not available as a meta data of tweet it will have to be computed. To do so Qwant will, upon request from the SocialTruth platform, retrieve a list of tweets and for the 10 firsts tweets of the results list will retrieve all their retweets and compute from their timestamp the retweet pace of the initial one. Special attention will be put on optimizing processing time for the whole process.

2.7 Fine grained Sentiment Analysis and Fake tweet detection

Recent work shown the interest to include sentiment analysis [Ajoa et al, 2019]. They show improvements in fake tweet detection by adding sentiment analysis features. This why we propose to add several sentiment analysis features, and more precisely, a standard polarity feature (negative, positive, neutral) [Saif et al., 2013] and we propose to go deeper by adding a fined grained sentiment analysis (strongly negative, weakly negative, neutral, weakly positive, strongly positive).

In this study we focused on an industrial way to implements sentiment analysis, since our internal requirements aim to process one query or sentence in 1 ms, we need to process a tweet using the same constrain. Against this constrain, very large approaches for sentiment analysis such as BERT [Munikar et al., 2019] or even lighter approach like CNN [Kim, 2014] are far too slow in our framework. To our knowledge, the most efficient and fastest approach is the FastText sentence classification [Joulin et al., 2014].

Our approach focuses on two points, the first one is in the data cleaning, in which we remove the usernames (beginning with “@”) and we remove the sign “#” from hashtag. We also keep emoticons and all signs which helps to detect sentiments in tweets. The model chosen is a simple layer classifier, using the FastText toolkit [Joulin et al., 2014] to stick our internal requirements.

We evaluated our selected approach on the SST-5 corpus [Socher et al., 2013] which is a well-known Fined-grained Sentiment Analysis benchmark. The dataset is composed a train and a test dataset, which respectively are composed of 8544 and 2210 sentences (163K and 44K words), using 5 cardinality, which has roughly balanced number of samples: strongly negative, weakly negative, neutral, weakly positive, strongly positive. Even if this corpus is structurally different, we can evaluate our approach and compare it to the Fined-Grained Sentiment Analysis state-of-the-Art.

Table 3 show the result of our approach. We trained our Word Embedding on all data available in SST-5. The approach we choose aims to use a sentence embeddings approach to classify sentences. In our context a tweet is considered as a sentence. We also reported several approaches mentioned previously, to show the distance against the State-of-the-Art.

Since the SST-5 corpus is well balanced, we used the accuracy metric, like in Munikar et al., [2019]. Our approach shows better results than simple approach using an average word embedding approach but is less efficient (41.6 points) than more elaborated approach such as RNTN [Socher et al., 2013], CNN [Kim, 2014] or BERT [Munikar et al., 2019] (respectively 45.7, 48.0, and 55.5). We need to conduct more research to enhance our approach to stick to the State-of-the-Art in a better way in spite of the time processing constrain.

Approaches	**Accuracy**
Average Word Vectors [Socher et al., 2013]	32.7
FastText (selected approach)	**41.6**
RNTN [Socher et al., 2013]	45.7
CNN [Kim, 2014]	48.0
BERT-Base [Munikar et al., 2019]	53.2
BERT-Large [Munikar et al., 2019]	**55.5**

Table 5: Results on STST-5 dataset.

3. Implementation

3.1 Assumptions

The collection and supply of the tweets will be done using the Twitter Developer API for (research context) that Qwant will pay for (using reallocation of its overall budget) as tweet collection is no more a part of day-to-day operation of the Qwant search engine. Provisioning of this service to the SocialTruth platform for the project's needs will be done through the Qwant Partners' API.

In the context of day-to-day operation, a maximum of 500 tweets will be provided in reply to a specific request from the SocialTruth platform via the Qwant API, but the Meta data "retweet pace" will be provided only for the 10 first of the list. Number of requests from the SocialTruth platform to the Qwant API will be limited to 1000 calls and 500 K Tweets per month during the project.

3.2 Requirements

To access the Qwant API; the SocialTruth platform needs access to Qwant server at: https://api.qwant.com/partners/v2. For each request the SocialTruth platform needs to send the Client ID and token provided by Qwant in each request.

3.3 Method and timeframe

The selected Qwant API (News or Tweets) accepts request using the GET method in Curl, as shown in example hereafter:

```
curl -X GET \
  'https://api.qwant.com/partners/v2/qwant/web?q=france' \
  -H 'Client-ID: xxx' \
  -H 'Token: xxx' \
```

3.4 Datasets constitutions for the purpose of training models

The use of "Social Mining" techniques to generate training datasets for SocialTruth models, is based on the hereabove implementation principles. But in order to build-up datasets valuable for fake news detection, Qwant will aim to collect a significant number of "fake" tweets.

Based on the corpus proposed by Shu et al.(2018), which contains 164,892 fake political posts, 399,237 factchecked political posts, 519,581 fake gossip posts and 876,967 factchecked gossip posts, we will extend the existing corpus by collecting new tweets, up to 2Millions of each kind of tweets (fakenews, factcheked, unknown).

4. Summary and conclusions

The social mining service that Qwant will provide to the SocialTruth platform via a Rest API, will allow the SocialTruth platform to get news and tweets enriched with sentiment and meta data in real time upon request via the API. These will then be analyzed within the SocialTruth platform and the Open Verification Ecosystem set within WP 4**.** The Rest API format that will be used will allow an easy integration within the platform in WP 5. This social mining system will be operational in month 24. Qwant will then take into account feedback from the consortium partners, and more specifically from ESF and UTP, to enhance its social mining system and more specifically its fine-grained sentiment analysis models.

5. References

Shu et. al, (2018) Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee and Huan Liu, "FakeNewsNet: A Data Repository with News Content, Social Context and Dynamic Information for Studying Fake News on Social Media", arXiv preprint, 2018

[Ajoa et al., 2019] AJAO, Oluwaseun, BHOWMIK, Deepayan, et ZARGARI, Shahrzad. "Sentiment aware fake news detection on online social networks". In : *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019. p. 2507-2511.

[Saif et al., 2013] Hassan SAIF, Miriam FERNANDEZ, Yulan HE, and Harith ALANI. "Evaluation datasets for twitter sentiment analysis: a survey and a new dataset", the sts-gold.2013

[Munikar et al., 2019] Manish MUNIKAR, Sushil SHAKYA, and Aakash SHRESTHA. "Fine-grained sentiment classification using BERT". In : 2019 Artificial Intelligence for Transforming Business and Society (AITB). IEEE, 2019. p. 1-5.

[Kim, 2014] Yoon KIM. "Convolutional Neural Networks for Sentence Classification". Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). 2014.

[Joulin et al., 2017] Armand JOULIN, Édouard GRAVE, Piotr BOJANOWSKI, and Thomas MIKOV. "Bag of Tricks for Efficient Text Classification". In : Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers. 2017. p. 427-431.

[Socher et al., 2013] Richard SOCHER, Alex PERELYGIN, Jean Y. WU, Jason CHUANG, Christopher D. MANNING, Andrew Y. NG and Christopher POTTS. "Recursive deep models for semantic compositionality over a sentiment treebank". In : Proceedings of the 2013 conference on empirical methods in natural language processing. 2013. p. 1631-1642.