

An Intelligent Dashboard for Assisted Tweet Composition in the Cultural Heritage Area (Work-in-progress)

Riccardo Martoglia

University of Modena and Reggio Emilia
Dip. Scienze Fisiche, Informatiche e Matematiche
Modena, Italy
riccardo.martoglia@unimore.it

Manuela Montangero

University of Modena and Reggio Emilia
Dip. Scienze Fisiche, Informatiche e Matematiche
Modena, Italy
manuela.montangero@unimore.it

ABSTRACT

Cultural Heritage institutions are nowadays using social media to communicate with citizens and tourists. However, providing actual effective communication is not an easy task, as every day millions of messages are posted through social media. Thus, getting visibility is not trivial. In this paper we present the architecture of a dashboard, accessible by mobile Android devices, to support museum social media managers in composing effective tweets by providing suggestions to improve message drafts. At this aim, the application exploits machine learning techniques over data related to tweets posted by museums in the past.

CCS CONCEPTS

• **Human-centered computing** → **Social media**; *Empirical studies in collaborative and social computing*;

KEYWORDS

Twitter, machine learning, dashboard, assisted composition, cultural heritage

ACM Reference Format:

Riccardo Martoglia and Manuela Montangero. 2020. An Intelligent Dashboard for Assisted Tweet Composition in the Cultural Heritage Area (Work-in-progress). In *6th EAI International Conference on Smart Objects and Technologies for Social Good (GoodTechs '20)*, September 14–16, 2020, Antwerp, Belgium. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3411170.3411272>

1 INTRODUCTION

In today globalized and ever-expanding market, companies invest countless resources to reach lasting visibility and to have strong impact on the online communities. However, writing successful tweets is tougher than most social media managers might think: aside from knowing when to post tweets for maximum exposure and visibility, it is necessary to know how to compose the tweet to generate the highest number of likes and retweets. For instance, it is crucial to determine which content features (images, multimedia content, etc.) perform best in order to help the tweet pop up on

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GoodTechs '20, September 14–16, 2020, Antwerp, Belgium

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7559-7/20/09.

<https://doi.org/10.1145/3411170.3411272>

social feeds [11]; the same holds for the usage of hashtags, mentions, links, and so on [5, 19].

In this paper, we concentrate on the Cultural Heritage (CH) domain and we present an ongoing work consisting in the creation of an innovative dashboard meant to assist museum media managers in writing successful tweets. The dashboard exploits intelligent data analytics and machine learning techniques to analyze past tweets history from the same museum and other similar ones and provides: (a) a mobile UI client, to help media managers compose and post impacting tweets; (b) intelligent server side features including: (b.1) *automatic tweet categorization* to predict the impact of the tweet, i.e., if the tweet is going to be successful or not written as it is; (b.2) *assisted tweet composition*, suggesting possible ways to improve the tweet and expect a better impact.

We measure the impact of a tweet by its number of likes. Indeed, placing a like is an intentional action taken by the reader and one might hope that those placing likes had at least read the tweet. The ultimate goal of this work in progress project [7–9] is to support effective communication and marketing campaigns in the CH domain.

The rest of the paper is organized as follows: in Section 2 we briefly report on related work; Section 3 gives an overview of the current project, while details are to be found in Sections 4 and 5. Finally, conclusions and future work are presented in Section 6.

2 RELATED WORKS

In the past few years, several studies focused on the analysis of social media data to promote social-based solutions for the CH sector. One of the main aims in this sector is on providing a fast and effective gathering of feedback from visitors, e.g., [3] provides a quantitative and qualitative approach to the analysis study of tweets posted during the MuseumWeek event organized by Twitter and [15] studies how the relation between archives, museums and users is changed thanks to the large diffusion of social networks.

Indeed, the use of social media data to understand user feelings and behaviour is not new: various approaches have been recently proposed in computer science literature for their analysis. For instance, a large body of work (e.g., [12, 16]) uses machine-learning [1] and/or lexicon-based [14] techniques to understand users' opinions, sentiments and attitudes towards different topics [2, 4, 6, 10]; other streams of research use social network analysis to identify social influencers [20]; others use psychological models to predict the users' psychological traits (see, for instance the appliedmagicsauce.com project developed at the University of Cambridge). However, there seems to be a lack of approaches exploiting social media data for proactively helping in composing more popular messages.

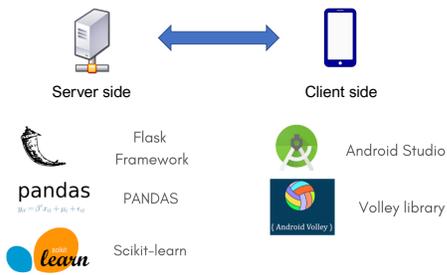


Figure 1: Dashboard architecture and technologies

The concept of dashboard as a user interface providing at-a-glance views relevant to a particular objective is a common one in the business sector. In particular, focusing on social media and Twitter, several examples exist. The Tweet activity dashboard¹ shows the user a detailed analysis of his/her activity on Twitter. For each sent Tweet, the user is able to track: the number of impressions, the number of engagements, the engagement rate (impressions divided by engagements). IOM’s Twitter dashboard² collates all the Organization’s global, regional, special liaison office, mission and special projects accounts. It provides a single repository for all the latest news about IOM’s work all over the world and in particular, latest updates on migration. All these examples, however, are not focused on the CH sector and do not exploit intelligent data analytics features to give useful suggestions to users.

3 DASHBOARD OVERVIEW

This project is developed for mobile devices, as it is meant to be easily accessible to a large public of users. The dashboard architecture consists of two main parts:

- a *server* application that stores all past tweet data and manages all “*intelligent*” tasks. In particular, the server handles a classifier model, receives incoming messages from the client and formulates suggestions in order to improve such messages. Classification is done depending on the expected “*success*” of the message in terms of number of likes;
- a simple and intuitive *client* application, working on Android devices providing UI and communication with the server.

The server side is written in Python and exploits the Pandas library for data manipulation and the scikit-learn library for machine learning and classification functionalities. The client side is written in Android studio and exploits the Volley library for networking.

The user performs the following tasks (see Figure 2):

- after login, the user enters his/her own message draft, and waits for the classifier evaluation (server side);
- the user gets a classification (either GOOD or BAD) for his/her draft message;
- in case of BAD outcome, the user can access and follow the suggestions given by the system in order to improve the impact of the message and new versions of the messages undergo classification again;

¹<https://business.twitter.com/it/analytics/tweet-activity-dashboard.html>
²<https://www.iom.int/dashboard/> (International Organization for Migration)

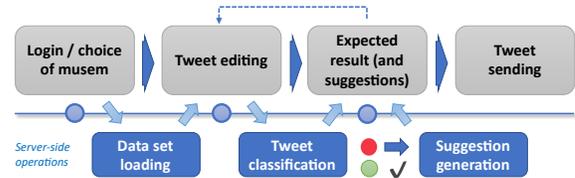


Figure 2: Dashboard usage flow

	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
	Museum of Modern Art	British Museum	Centre Pompidou	MFA Boston	Visit MUVE	Egyptian Museum
	Tate	Van Gogh Museum	National Gallery	Museums in common	Picasso Museum (Barcelona)	Uffizi
	MET Museum	Louvre Museum	Frida Kahlo Museum	Phila Museum	Mart Museum	MUSE Trento
	Guggenheim	Getty Museum	Orsay Museum	Masp museum		
	Saatchi Gallery	Prado Museum		National Gallery of Art		
				MAXXI Museum		
Total Tweets	3761	3838	2474	4931	2164	2420

Table 1: Dashboard architecture and technologies

- once the user is satisfied with the outcome, he/she can log on to Twitter and post the message.

In the next sections will give details of the server and client sides.

4 SERVER SIDE

Data preprocessing. To set up the server side, we considered 26 well known world spread art museums and partitioned them according to the number of followers of their Twitter account, so that they are grouped by popularity representing a wide range of situations. We got six groups that are listed in Table 1. We analyzed ca. 20k tweets, with a number in within 2k and 5k for each group. The user messages will be handled by our system by taking into consideration only the tweets of the museums in the group more representative of the user’s museum, i.e., the group with the number of followers closest to the number of followers of the user museum (or the group in which the museum is in, if the user represents one of the 26 chosen museums).

We used the Pandas library to build a DataFrame containing all information we needed for our system to work. Each row of the table represents a tweet and contains the following features: NLIKE (the amount of likes received), NURLS (number of url links contained in the message), NIMG (amount of images in the tweet), NHASH (number of hashtags contained), NMENTION (number of mentions contained), LENGTH (length of the message in chars), SENT (a Boolean value that is used to indicate the overall sentiment of the message, i.e. positive or negative, as derived from sentiment analysis).

Features	Accuracy %					
	Group					
	1	2	3	4	5	6
NURLS - NIMG - NHASH - NMENTION - LENGTH - SENT	0,79	0,8	0,78	0,80	0,80	0,79
NIMG - NHASH - NMENTION	0,77	0,78	0,77	0,79	0,78	0,80

Table 2: Accuracy for different groups and feature sets

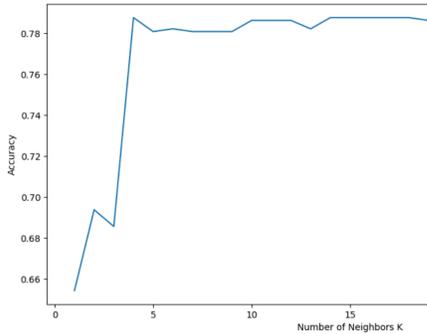


Figure 3: Choice of optimal K for KNN classifier

Automatic tweet classification. We defined a simple binary tweet classification into GOOD and BAD ones: a tweet is good or bad depending on the number of likes it got, group by group (i.e., NLIKE was our target variable). Intuitively, if the number of likes of a given tweet of a museum in a given group is close to the maximum number of likes received by the tweets of the same group, then the tweet is classified as GOOD, dual for BAD ones.

The classification model exploits the KNN algorithm from the Scikit-learn library; the choice of the model, based on a similarity between the tweet features, is novel w.r.t. our past works [7–9] and is also motivated by the way suggestions are generated (more on this later). We employed the standard euclidean distance as distance metric between the features. Features are normalized and scaled in the range [0,1] after each test-validation-train split.

We performed a large number of tests in order to understand, for each group, the impact of the different features on the classification and to assess the overall classification accuracy. All the results are the average of 10 runs of execution. Table 2 shows the accuracy achieved for two different feature sets: the whole set of features, and a selection of the three most significant features detected by our tests (NIMG, NHASH, NMENTION). Notice that the accuracy in the two cases are very near, therefore we decided that working on a reduced amount of features was a good thread off between model complexity and accuracy level. Moreover, we determined test-validation-training split ratios (10% testing, 10% validation and 80% training); through validation we determined the value K for the KNN classifier. In particular, for each validation data set, (i.e., for each group of museums) we performed a set of tests where we repeatedly assigned increasing K values (from 1 to 20) and then averaged the obtained accuracy values. Figure 3 shows the obtained accuracy trend w.r.t. K : as we can see, accuracy increases and then stabilizes: we chose $K = 5$ as the best (i.e., minimum odd) value for our configuration.

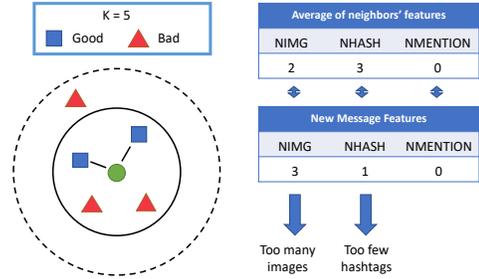


Figure 4: Suggestion generation

Suggestion generation. Our classification framework based on a simple distance metric and few numerical features, allows us to identify nearby “positive example” tweets for a given message t , and formulate feature-based suggestions to improve such message. In this initial phase of the project, we consider simple suggestions such as increasing/decreasing the value of the involved features (e.g., add/remove an image or a mention); we plan to consider more advanced suggestions in the future. Suggestion generation is triggered whenever a draft message t receives a negative (BAD) classification, and works following these steps:

- (1) *KNN search*: the K nearest neighbors to the original tweet t are extracted by means of a KNN search on the dataset of the proper museum group and denoted by t_{nn} ;
- (2) *Negative neighbor filtering*: all neighbors in t_{nn} that are categorized as BAD are removed from t_{nn} , as they would give no useful suggestion;
- (3) *Distant neighbor filtering*: we keep only those tweets in t_{nn} that are closer to t . Indeed, even if such tweets belong to the same cluster of t , their distance might still be large: messages with very different characteristics from the initial one would be difficult to “mimic” without requiring a complete rewrite. Therefore, we define a threshold T and drop the tweets in t_{nn} whose distance from t exceeds T ;
- (4) *Computation of improvement feature values*: “improvement” feature values are computed by considering an aggregate of the features of the tweets in t_{nn} . In particular, we choose to simply compute the mathematical average of their numerical values, as we are interested in the general trend and in dampening pronounced differences among the values;
- (5) *Suggestion generation*: finally, improvement features are compared to the features of the message t and suggestions are provided to make the latter closer to the former.

Note that, for better consistency between classification and suggestion generation, the considered features and value of K are the same as the ones used for classification (i.e., number of images, hashtags and mentions and $K = 4$); also, the same Scikit-learn library is exploited for nearest neighbor search. Furthermore, in case of an empty t_{nn} at the end of the filtering steps, suggestion generation terminates and no suggestions are returned, since no useful suggestions could be extracted from the current dataset (this seldom happened in our experiments).

An example of the suggestion generation process is given in Figure 4, in which t is the small green circle. In step (1) 4 tweets

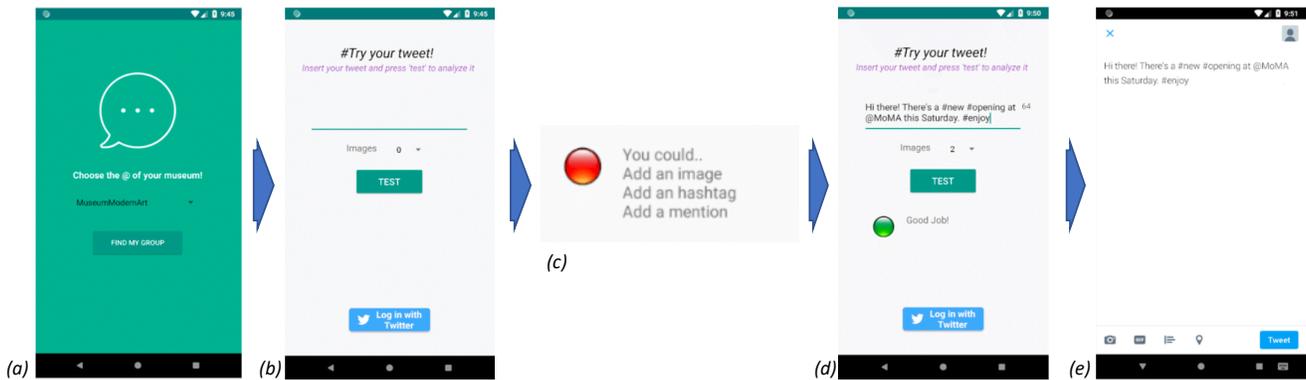


Figure 5: Dashboard UI and usage example

nearest to t are identified: three of them are positive (blue squares) and one is negative (red triangle). The negative one is dropped in step (2). Step (3) filters out a positive message that is beyond the threshold T (depicted as a solid circle). Target features (step 4) and generated suggestions (step 5) are shown on the right side: being the number of images of t (3) higher than the improved one (2) a suggestion such as “You should drop an image” is generated; the opposite happens for hashtags (“You should add an hashtag”).

5 CLIENT SIDE

The current prototype of the dashboard UI (Figure 5), created in Android Studio, is very basic and still preliminary, anyway it already supports the main operations described in Section 3. The dashboard is designed to be simple and intuitive, in order to be accessible to a multiplicity of users with different technical skills. Through the login / initial selection screen (Figure 5-a), the user can choose his own museum (or select a group of similarly sized museums). The list of museums is dynamically filled at the moment that the program starts on the basis of the data available in the server. The main screen (Figure 5-b) is used to edit the draft of the tweet and to receive composition assistance. Once the user has entered the draft, this is sent to the server that starts the analysis process and shows (in case of negative classification) possible suggestions (Figure 5-c). Then, the draft can be edited accordingly and, when the user is satisfied (Figure 5-d), he/she can login to Twitter and finally tweet the message (Figure 5-e).

6 CONCLUSIONS AND FUTURE WORK

In this work-in-progress paper we have presented the design of an intelligent dashboard to assist museum social media managers in composing successful tweets. In our wider on going project we plan to support effective communication and marketing campaigns for the CH domain: we plan to consider additional ways of generating suggestions (e.g., adapting semantic customization techniques successfully tested in other scenarios [13, 17, 18]) and to extend the set of possible suggestions (e.g., additional features as the time of the day when to post the tweet, specific words in the tweet text and so on).

ACKNOWLEDGMENTS

We thank Chiara Boni for her work in programming the dashboard.

REFERENCES

- [1] N. Aston, J. Liddle, and W. Hu. 2014. Twitter Sentiment in Data Streams with Perceptron. *Journal of Computer and Communications* 02 (01 2014), 11–16.
- [2] A. Bujari, M. Furini, and N. Laina. 2017. On using cashtags to predict companies stock trends. In *Proc. CCNC 2017*. IEEE, 25–28.
- [3] A. Courtin, B. Juanals, J. L. Minel, and M. d. S. Léger. 2014. The Museum Week Event: Analyzing Social Network Interactions in Cultural Fields. In *Proc. of the 10th Int. SITIS Conf.* 462–468.
- [4] R. De Michele and M. Furini. 2017. Smart City and Images: The Use of Image Hashtags to Get Insights on Citizens. In *GOODTECHS 2017*, Vol. 233. Springer, 311–320.
- [5] R. De Michele and M. Furini. 2017. The Use of Hashtags in TV-Shows: Analysis and Guidelines. In *Proc. GOODTECHS 2017*, Vol. 233. Springer, 321–330.
- [6] S. Ferretti, M. Furini, and M. Montangelo. 2019. Diabetes: What Are Italian Twitter Users Talking About?. In *Proc. ICCCN 2019*. IEEE, 1–6.
- [7] M. Furini, F. Mandreoli, R. Martoglia, and M. Montangelo. 2017. The Use of Hashtags in the Promotion of Art Exhibitions. In *Proc. 13th Italian Research Conference on Digital Libraries (IRCDL), Revised Selected Papers*. 187–198.
- [8] M. Furini, F. Mandreoli, R. Martoglia, and M. Montangelo. 2018. Towards Tweet Content Suggestions for Museum Media Managers. In *Proc. of GOODTECHS18*. Association for Computing Machinery, 265–270.
- [9] M. Furini, M. Mandreoli, R. Martoglia, and M. Montangelo. 2018. 5 Steps to Make Art Museums Tweet Influentially. In *Proc. 3rd Int. Workshop on Social Sensing*.
- [10] M. Furini and M. Montangelo. 2018. Sentiment analysis and Twitter: a game proposal. *Pers. Ubiquitous Comput.* 22, 4 (2018), 771–785.
- [11] M. Furini and M. Montangelo. 2019. On Predicting the Success of Political Tweets Using Psycho-Linguistic Categories. In *Proc. ICCCN*. IEEE, 1–6.
- [12] A. Giachanou and F. Crestani. 2016. Like It or Not: A Survey of Twitter Sentiment Analysis Methods. *ACM Comput. Surv.* 49, 2 (2016), 28:1–28:41.
- [13] F. Grandi, F. Mandreoli, and R. Martoglia. 2009. Issues in Personalized Access to Multi-Version XML Documents. In *Open and Novel Issues in XML Database Applications: Future Directions and Advanced Technologies, IGI Global*. 199–230.
- [14] X. Hu, J. Tang, H. Gao, and H. Liu. 2013. Unsupervised Sentiment Analysis with Emotional Signals. In *Proc. of the 22nd Int. Conf. on World Wide Web*. 607–618.
- [15] B. Jensen. 2013. Instagram as cultural heritage: User participation, historical documentation, and curating in Museums and archives through social media, Vol. 2. 311–314.
- [16] S. E. Kase and E. K. Bowman. 2018. Operating in the New Information Environment: An Army Vision of Social Sensing?. In *2018 International Workshop on Social Sensing (SocialSens)*. IEEE Computer Society, Los Alamitos, CA, USA, 1–1.
- [17] F. Mandreoli, R. Martoglia, P. Tiberio, et al. 2006. Semantic Web Techniques for Personalization of eGovernment Services. In *Proc. ER SemWAT 2006*. 435–444.
- [18] R. Martoglia. 2011. Facilitate IT-Providing SMEs in Software Development: a Semantic Helper for Filtering and Searching Knowledge. In *Proc. of SEKE*. 130–136.
- [19] R. De Michele, S. Ferretti, and M. Furini. 2019. On helping broadcasters to promote TV-shows through hashtags. *Multim. Tools Appl.* 78, 3 (2019), 3279–3296.
- [20] J. Yang and J. Leskovec. 2010. Modeling Information Diffusion in Implicit Networks. In *Proc. of ICDM*. IEEE Computer Society, 599–608.