

Is Readability a Valuable Signal for Hashtag Recommendations?

Ion Madrazo Azpiazu
Computer Science Dept.
Boise State University
Boise, Idaho, USA
ionmadrazo@boisestate.edu

Maria Soledad Pera
Computer Science Dept.
Boise State University
Boise, Idaho, USA
solepera@boisestate.edu

ABSTRACT

We present an initial study examining the benefits of incorporating readability indicators in social network-related tasks. In order to do so, we introduce TweetRead, a readability assessment tool specifically designed for Twitter and use it to inform the hashtag prediction process, highlighting the importance of a readability signal in recommendation tasks.

CCS Concepts

•Human-centered computing → Social recommendation; Social networks;

Keywords

Hashtag Recommendation; Readability

1. INTRODUCTION

Readability is a measure of the ease with which a text can be read. Usually represented by a number, it is an indicator used by teachers to classify and find appropriate resources for students. Several studies have demonstrated the benefits of using readability indicators in educational-related applications, such as book recommendation, text simplification, or automatic translation. However, applying readability indicators outside this environment remains relatively unexplored. Social networks could benefit from readability assessment. Twitter is a social network where users and texts are the main focus. For this reason, it is natural to think that for Twitter the ease with which a tweet can be understood by a user may affect his interest in it, and therefore influence actions taken, such as re-tweeting, giving a like or replying to the tweet.

The authors of [6] examined the degree to which the age of a user, a feature strongly correlated with readability, influences who people follow on Twitter, and demonstrated that Twitter users have a higher chance to follow people of similar age. Using standard readability measures in text from Twitter, which constrains tweets to be of at most 140 characters in length, is not a trivial task. The lack of structure and shortness of those texts make standard natural

language analysis techniques inefficient. With that in mind, we developed TweetRead, a novel readability assessment tool specifically designed for tweets. TweetRead takes advantage of social information, such as hashtags or mentions, for predicting the text complexity levels of tweets. Furthermore, in order to highlight the usefulness of such a tool in social networking environments, we developed a simple, yet effective, hashtag recommendation strategy that takes advantage of TweetRead-generated complexity levels of tweets to inform the hashtag recommendation process.

2. TWEETREAD

TweetRead’s goal is to estimate readability of any given tweet T . TweetRead is based on a logistic regression technique¹ that fuses simple indicators describing T from different perspectives and determines its text complexity. The indicators considered by TweetRead include: (i) T ’s readability level, estimated using *Flesch*² [1], (ii) T ’s similarity with respect to word distributions generated from a large Twitter corpora C labeled by age groups, (iii) average readability of each hashtag h in T , computed based on the average readability levels estimated using *Flesch* of tweets in C that include h , (iv) average readability level of users mentioned on T , estimated using *Flesch* on tweets written by mentioned users, and (v) frequency of mentions, emoticons, and hashtags in T .

Unlike traditional readability formulas that tend to map readability levels with school grades, to tailor TweetRead to the Twittersphere, we consider six levels of text complexity following Levinston’s [3] adult development stages.

3. HASHTAG RECOMMENDATION

Hashtags are character strings used to represent concepts on Twitter, starting with a # symbol. They are a core Twitter feature and serve classification and search purposes. Their unrestricted nature, however, creates difficulties, including the fact that the same concept can be represented by different hashtags, hindering the search process of a concept [5]. For example, tweets related to the Monaco Formula 1 Grand Prix can be searched using #monacoGP, #monacoF1GP or #monacoF1 retrieving different results. Hashtag recommendation aims at identifying suitable hashtags a user can include in his tweet to reduce the space of tags generated [5] and facilitate the ease with which he and other users can locate the corresponding tweet.

Given that (i) the scope of this paper is to validate the importance of considering a text complexity signal to enhance

¹We empirically verified that among numerous supervised techniques, logistic regression was the most promising one.

²Flesch estimates the readability of a text/tweet t , by examining its length and the average length of terms in t .

a recommendation task and (ii) multiple and increasingly complex systems have been developed for hashtag recommendation [2], we base our study on an existing framework for hashtag recommendation presented in [5]. Given a tweet T , the proposed framework identifies existing hashtags to recommend by following two major steps: (1) generate candidate hashtags by recommending hashtags present in similar tweets, using tf-idf based cosine similarity and (2) rank hashtags from retrieved candidate tweets using different strategies. The strategies presented in [5] include:

- **Similarity.** Prioritizes hashtags included on tweets that have the closes similarity to T , as estimated using the well-known tf-idf and cosine similarity measure.
- **Global popularity.** Prioritizes hashtags based on their respective frequency of occurrence on Twitter.
- **Local popularity.** Prioritizes hashtags based on their frequencies of occurrence among the tweets retrieved in response to T .

We enhance the proposed strategies by taking advantage of TweetRead, as follows:

- **TweetRead.** Prioritizes candidate hashtags that have the same or similar text complexity (estimated using TweetRead) with respect to T .
- **PopularityTweetRead.** Prioritizes hashtags based on their frequencies of occurrence among tweets whose readability level is estimated to match T 's.
- **SimilarityTweetRead.** Prioritizes candidate hashtags based on their respective ranking scores computed using Similarity only on tweets whose readability level is estimated to match T 's .

4. INITIAL ASSESSMENT

In this section, we discuss an initial evaluation on TweetRead, as well as its applicability for suggesting hashtags.

TweetRead. Given that readability of social content is an unexplored area, benchmark datasets that can be used for evaluation purposes are unavailable. For this reason, we built our own dataset. We initially gathered 172M tweets over an 8-month period using Twitter streaming API. For the purpose of this experiment we assume that the age of people exactly corresponds to their readability level, and that each tweet written by a user will have the same readability level as its author. With that in mind, we followed the framework presented in [6], which examines patterns such as “happy xth birthday”, for determining the age of Twitter users. In doing so, we eliminated from our dataset, users (and their corresponding tweets) from whom age could not be determined. Thereafter, we grouped labeled tweets into 6 age groups, which translates into a uniformly distributed dataset of 22k tweets with their corresponding readability levels. We followed a 10-cross-fold validation strategy and measured the accuracy of the predicted readability levels with respect to the ground truth. As shown in Table 1, TweetRead significantly outperforms the baselines considered for this assessment: Flesch [1] and Spache [4], which are two well-known, traditional readability measures. The reported results demonstrate the need for readability strategies that examine information beyond standard text analysis, if they are meant to be successfully used in the social networking context.

Hashtag recommendation. For evaluating the strategies for hashtag recommendation presented in Section 3, we

Flesch	Spache	TweetRead
27%	31%	81%

Table 1: Performance evaluation of TweetRead vs. baselines.

used the aforementioned dataset. We treated the hashtag of each corresponding tweet as the ground truth. In other words, for each tweet T , we generated the corresponding top-N hashtag recommendations and considered relevant the ones matching the hashtags in T . As in [5], we used the recall measure to evaluate performance and determine to which extend the correct hashtags were recommended within the top N generated suggestions. As shown in Figure 1, even if readability on its own is not a sufficient factor to suggest hashtags, when combined in-tandem with other content-based and/or popularity strategies, it leads to the improvement of the overall hashtag recommendation process.

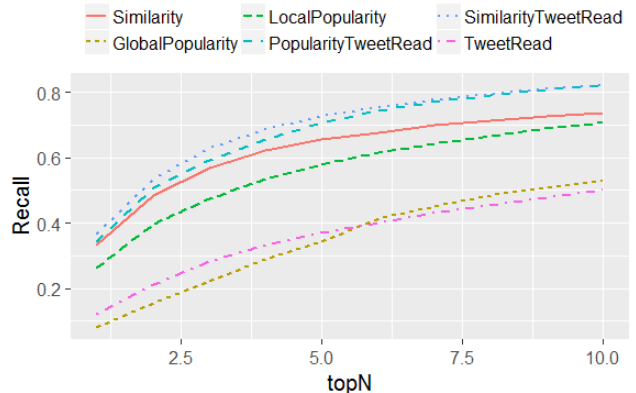


Figure 1: Hashtag recommendation assessment.

5. CONCLUSION AND FUTURE WORK

In this paper, we presented TweetRead, a novel readability assessment tool specifically designed to predict the readability of tweets. We also discussed the initial study conducted to demonstrate the benefit of using a readability signal in the hashtag recommendation task, which yielded promising results. In the future, we plan to explore other applications of readability in social networks, such as user recommendation, advertisement targeting or re-tweet prediction. We will also explore techniques to further enhance TweetRead and adapt it to other social networks beyond Twitter.

6. REFERENCES

- [1] R. Flesch. A new readability yardstick. *Journal of Applied Psychology*, 32(3):221, 1948.
- [2] F. Godin, V. Slavkovicj, W. De Neve, B. Schrauwen, and R. Van de Walle. Using topic models for twitter hashtag recommendation. In *WWW*, pages 593–596. ACM, 2013.
- [3] D. J. Levinson. A conception of adult development. *American psychologist*, 41(1):3, 1986.
- [4] G. Spache. A new readability formula for primary-grade reading materials. *The Elementary School Journal*, 53(7):410–413, 1953.
- [5] E. Zangerle, W. Gassler, and G. Specht. Recommending #-tags in twitter. In *SASWeb 2011*, volume 730, pages 67–78, 2011.
- [6] J. Zhang, X. Hu, Y. Zhang, and H. Liu. Your age is no secret: Inferring microbloggers’ ages via content and interaction analysis. In *AAAI ICWSM*, 2016.