

Emotional Social Signals for Search Ranking

Ismail Badache
LSIS UMR 7296 CNRS
University Aix-Marseille
Marseille, France
Ismail.Badache@lsis.org

Mohand Boughanem
IRIT UMR 5505 CNRS
University Paul Sabatier
Toulouse, France
Mohand.Boughanem@irit.fr

ABSTRACT

A large amount of social feedback expressed by social signals (e.g. like, +1, rating) are assigned to web resources. These signals are often exploited as additional sources of evidence in search engines. Our objective in this paper is to study the impact of the new social signals, called Facebook *reactions* (*love, haha, angry, wow, sad*) in the retrieval. These reactions allow users to express more nuanced emotions compared to classic signals (e.g. like, share). First, we analyze these reactions and show how users use these signals to interact with posts. Second, we evaluate the impact of each such reaction in the retrieval, by comparing them to both the textual model without social features and the first classical signal (*like*-based model). These social features are modeled as document prior and are integrated into a language model. We conducted a series of experiments on IMDb dataset. Our findings reveal that incorporating social features is a promising approach for improving the retrieval ranking performance.

CCS CONCEPTS

• **Information systems** → *Document filtering*;

KEYWORDS

Facebook Reactions; Social Signals; Social IR; Ranking

1 INTRODUCTION

Majority of information retrieval (IR) systems exploit two classes of features to rank documents in response to the user query. The first class, the most exploited, depends on the query, it concerns term statistics such as term frequency, distribution of term in documents. The second class concerns query-independent features, which measures a kind of quality or importance of the document. Among these factors, the number of incoming links to a document [10], PageRank [5], topical locality [7], the presence of URL [13], document authors [11] and social signals [1–3].

Most of existing approaches [2, 3, 6] exploit non-emotional signals such as (+1, *share, tweet*) to estimate the document prior by considering the quantity of signals related to a resource. In this paper we are interested in a novel type of signals, named Facebook

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.
SIGIR '17, August 07–11, 2017, Shinjuku, Tokyo, Japan.

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5022-8/17/08...\$15.00

<https://doi.org/10.1145/3077136.3080718>

reactions. We assume that, these emotional reactions (*love, haha, wow, angry, sad*) associated to a web resource (document) can be seen as clues that indicate a resource interest beyond a social network or a community. The research questions are the following:

- (1) How users use these reactions to interact with resources?
- (2) What is the impact of Facebook reactions on IR?

The remainder of this paper is organized as follows: Section 2 reviews some background and related work. Section 3 presents a statistical analysis on Facebook reactions. Section 4 describes our social approach. Section 5 reports on the results of our experimental evaluation. Finally, Section 5 concludes this paper with some perspectives.

2 BACKGROUND AND RELATED WORK

This section reports: (1) some background information about social signals and; (2) related work that has leveraged social signals to measure a priori relevance of a resource.

2.1 Social Signals

Social signals represent one of the most popular UGC (User Generated Content) on the Web. Indeed, the Web pages include buttons of different social networks where users can express whether they support, recommend or dislike content (text, image, video, etc). These buttons which describe social activities' actions (e.g., *like, share, +1*, etc) are related to specific social networks (e.g., Facebook, Google+, etc) with counters indicating the rate of interaction with the Web resource. In February 2016, Facebook introduces additional *emotional* signals (reactions), allowing users to interact with posts (resources) across *love, haha, wow, angry, and sad* (see Figure 1). These reactions are an extension of the *like* button, to give users more ways to express their feelings towards a post in a quick and easy way. The goal of these new signals is to encourage users to react even if the contents are difficult to *like* as in the case of disasters, gloomy news, death, emotion on movie. Table 1 summarizes the most popular signals on social networks.

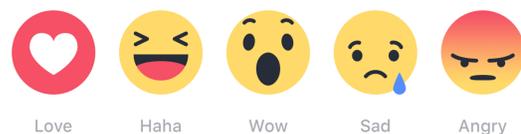


Figure 1: Additional Facebook signals (reactions)

2.2 Exploiting Social Signals in a Search

Social signal assigned to a given resource can be interpreted as an approval or disapproval of a resource, which can be used to measure a rank of importance of a resource.

Table 1: List of different social signals types

Type	Example	Social network
Vote	Like +1	Facebook, LinkedIn, Google+, StumbleUpon
Message	Tweet Post	Facebook, Google+, LinkedIn, Twitter
Share	Share Re-tweet	Google+, Twitter, Buffer, Facebook, LinkedIn
Tag	Bookmark Pin	Delicious, Diigo, Digg Pinterest
Comment	Comment Reply	Facebook, Google+, LinkedIn, Twitter
Emotion	Love, Haha, Wow Sad, Angry	Facebook
Relation	Followers Friends	Facebook, Twitter

While several works have been exploiting users feedback there is still lack of studies that would analyze users' signals coming from specific social networks. Major existing works [2, 3, 6, 9, 12, 14] focus on how to improve IR by exploiting users' actions and their underlying social network. The most related work to ours is [12], which leveraged user feedback about YouTube videos for the task of affective video ranking. The videos are represented with sentic-features (based on the primary emotions) as well as other social features, and the queries are again assumed to include an explicit emotion statement. Learning to rank approach is used, which allows to compare the performance of different sets of features.

Our work has a similar motivation as those previous efforts, i.e., harnessing social features around a Web document to improve relevance ranking of conventional text search. However, our approach attempts to exploit and evaluate the impact of emotional signals in the retrieval ranking.

3 FACEBOOK REACTIONS ANALYSIS

In this section, we conducted a preliminary study on how users use these reactions to interact with posts. We studied the reactions on 10 IMDb Top Box Office movies released in 2016 as well as 4650 articles published on some famous International Media between March 2nd and June 2nd, 2016: "The Guardian (UK)", "The Wall Street Journal (USA)", "The New York Times (USA)", "The Washington Post (USA)", "China Daily (China)", "The Times of India (India)", "The Sydney Morning Herald (Australia)", "Fox News (USA)", "Dawn (Pakistan)", "CNN (USA)". The number of each reaction for each post related to media or movies is collected using Facebook API and parsing Facebook pages.

Table 2: Reactions audiences in International Media

Media	Love	Haha	Wow	Angry	Sad
Washington Post	63%	12%	11%	8%	6%
New York Time	61%	13%	12%	8%	6%
Fox News	60%	14%	10%	8%	8%
Guardian	52%	9%	3%	25%	11%
CNN	50%	10%	28%	10%	2%
China Daily	48%	2%	10%	2%	38%
Wall Street Journal	24%	14%	25%	15%	22%
Sydney Morning Herald	21%	3%	18%	49%	9%
Times of India	11%	6%	3%	40%	40%
The Dawn	10%	18%	18%	40%	14%
Total	40%	10%	14%	20%	16%

We looked at which posts users used the *love*, *haha*, *angry*, *wow*, and *sad* reactions the most, as a percentage of all reactions. Tables 2 and 3 shows the percentage distribution of each reaction on both datasets: international media posts and IMDb Top Box Office posts.

Table 3: Reactions audiences in IMDb Box Office

Media	Love	Haha	Wow	Angry	Sad
Captain America	91%	9%	0%	0%	0%
The Jungle Book	88%	8%	2%	0%	2%
Conjuring 2	87%	1%	7%	2%	3%
Warcraft	81%	3%	14%	2%	0%
Me Before You	80%	7%	3%	0%	10%
X-Men: Apocalypse	78%	13%	9%	0%	0%
Now You See Me 2	65%	20%	13%	1%	1%
Alice in Wonderland 2	56%	30%	11%	2%	1%
Ninja Turtles 2	55%	33%	10%	1%	1%
Angry Birds	52%	40%	3%	3%	2%
Total	73%	16%	7%	1%	2%

Overall in tables 2 and 3, we found that users in the world use *love* reaction (40% for international media and 73% for IMDb Top Box Office) of all reactions. Users were less likely to use the two negative reactions (*angry* and *sad*) in movies, and funny reactions (*haha* and *wow*) in media.

We also analyzed Facebook users reactions to the first breaking news report about the March 22 Brussels attack. In the early morning of March 22, there were explosions in Brussels near the Zaventem airport and a subway car. The first posts about the attack were very similar across several media. We have listed 4 news published on official facebook pages of 4 media as well as the headline of the article that they linked to in their posts (see Table 4). Although the posts were made early morning, they received an average of 4200 total reactions. We found that with this highly emotional tragedy, users were more likely to use reactions other than *like* signal. In fact, the "CNN" and the "Fox News" followers used *nonlike* reactions more than they liked.

Table 4: List of articles (Brussels attack) published on official Facebook pages of 4 international media

Media	News Title
CNN	Breaking: Belgium police at the airport at Zaventem told CNN that 'there has been an explosion' and 'something has happened'.
Fox News	Breaking News: At least 13 Killed in explosions at Brussels airport, metro station, Belgian media report.
Sydney Morning Herald	Multiple explosions reported in Brussels, Belgium.
Times of India	Brussels airport explosion: Several feared dead, shouts in Arabic.

According to Figure 2 that investigates reactions individually, "CNN" and "Sydney Morning Herald" have the angriest following with 71% and 73%, respectively. The "Times of India" had the saddest following in response to the article post about the attack with 68%, and "Fox News" followers were between the sadness (49%) and the anger. Although there were some *wow* reactions, and a few scattered *love* and *haha* reactions, (thankfully) they accounted for fewer than one ten-thousandth of a percent of global reactions.

Finally, the use of reactions by Facebook users, we could have the ability to measure sentiment in an even more dynamic way. The *like* button already allows to know the contents that interested the user. Now, with additional five *emoji*, we can record more nuanced reactions. Through these reactions, the user provides information about what he *loves* or what he gets *angry*. By considering all this data, we can draw up a detailed profile of tastes and personality of the user. Our study of these social reactions is only to understand how people use them. However, these new signals can be exploited in sentiment analysis or considered as a priori knowledge to estimate the relevance of the document in response to a user information needs (see next section).

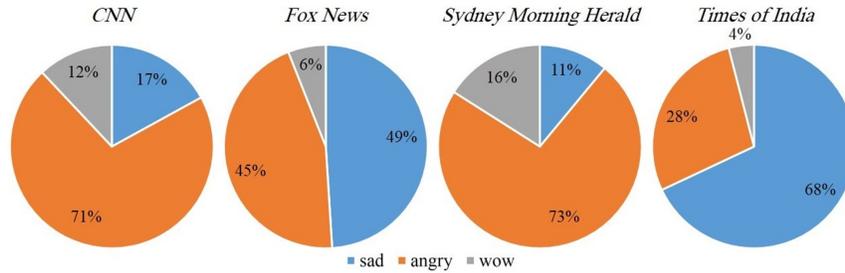


Figure 2: Media followers reactions to March 22 Brussels Attack

4 IMPACT OF REACTIONS ON IR

Our approach consists of exploiting Facebook reactions as prior knowledge to take into account in the search ranking. Textual relevance of a document is combined with its social importance estimated through these users reactions.

Social information exploited within the framework of our model is represented by $\langle U, C, R \rangle$ where $U=\{u_1, u_2, \dots, u_h\}$, $C=\{D_1, D_2, \dots, D_n\}$ and $R=\{r_1, r_2, \dots, r_m\}$ are finite sets of instances: *Users*, *Documents* and *Reactions*, respectively.

4.1 Textual Relevance and Document Priors

Textual model is used to estimate the relevance of a document to a query. Our approach combines the social document prior $P(D)$ and the relevance status value $RSV_{textual}(Q, D)$ between query Q and document D as follows:

$$RSV(D, Q) \stackrel{\text{rank}}{=} P(D) \cdot RSV_{textual}(Q, D) \quad (1)$$

$$\stackrel{\text{rank}}{=} P(D) \cdot \prod_{w_i \in Q} RSV_{textual}(w_i, D), \quad (2)$$

where w_i represents term in the query Q and $RSV_{textual}$ can be estimated with different models such as BM25 and language model. The document prior $P(D)$ is useful for representing and incorporating other sources of evidence to the retrieval process. Our main contribution is a method to estimate $P(D)$ by exploiting Facebook Reactions.

4.2 Estimating Priors

The priors are estimated by simply counting the number of reactions performed on the documents. Assuming that these features are independent, the general formula for $P(D)$ is:

$$P(D) = \prod_{r_i \in R} P(r_i), \quad (3)$$

where $P(r_i)$ is estimated using maximum-likelihood:

$$P(r_i) = \frac{|r_i(D)|}{|r_{\bullet}(D)|}, \quad (4)$$

where $|r_i(D)|$ is the number of reactions of type r_i on document D and $|r_{\bullet}(D)|$ is the total number of reactions on document D . Further, we use Dirichlet to smooth $P(r_i)$ by collection C to avoid zero probabilities. This leads to:

$$P(D) = \prod_{r_i \in R} \left(\frac{|r_i(D)| + \mu \cdot P(r_i|C)}{|r_{\bullet}(D)| + \mu} \right), \quad (5)$$

where $P(r_i|C)$, by analogously to $P(r_i)$, is estimated using maximum-likelihood:

$$P(r_i|C) = \frac{\sum_{D \in C} |r_i(D)|}{\sum_{D \in C} |r_{\bullet}(D)|} \quad (6)$$

Our approach can be useful for the queries that explicitly state a certain emotion (like “romance movies by...”, “happy videos of cats”). Indeed, using fine-grain emotional signals rather than a simple count of *like*, can improve IR where the user’s information need is related to her affective needs.

5 EXPERIMENTAL EVALUATION

To validate our approach, i.e. evaluate the impact of Facebook reactions on IR, we conducted a series of experiments on INEX IMDb¹ dataset (167,438 documents, 30 topics and qrels). Each document describes a movie, and is represented by a set of metadata, and has been indexed using keywords extracted from textual fields such as (*title*, *plot*, *actors*, etc). For each movie, Facebook reactions are collected via its API using movies URLs (see table 5). The nature of these signals is a counting of each reactions on movies. In our study, we focused on the top 1000 results for each topic.

Table 5: Instance of document with social reactions

Film Title	Facebook					
	Like	Love	Haha	Wow	Sad	Angry
Sinister	14763	8520	12	10256	647	146

We compared our approach which combines social document prior with Hiemstra language model [8], to the baselines *LM.Hiemstra* without using any document priors, as well as *like*-based model (*like* as document prior). We note that the best value of the smoothing parameter $\mu \in [90, 100]$.

5.1 Results and Discussion

Table 6 compares the different configurations of our approach in terms of precision@ k ($k \in \{10, 20\}$), nDCG@1000 and MAP. Different configurations are evaluated, by taking into account Facebook reactions individually and grouped according to their meaning: positive or negative emotions (labeled + or - in table 6, respectively). Label (*) indicates a statistically significant improvement over *LM.Hiemstra* with the significance level of 0.05.

a) Impact of Reactions Individually. The best results are obtained by the reactions *Love* and *Wow* with 0.4122 and 0.4031 in terms of P@10, respectively. Some topics such as “romance movies by Leonardo DiCaprio or Tom Cruise” and “romance movies by Richard Gere or George Clooney” are registered the highest precision when the reaction *Love* is taken into account in the ranking process (with 0.8801 and 0.9112 in terms of P@10, respectively). The reaction *Wow* is more effective with topics that represents a real fact or a freaky and exciting information such as “true story event movies” that registered a precision at ten of 0.7754. The reactions

¹<https://inex.mmci.uni-saarland.de/tracks/dc/2011/>

Haha and *Sad* provide the lowest results compared to *like*-based model, but they bring significant improvements compared to textual model *LM.Hiemstra*. The *Angry* is the weakest feature, it is close to a negative signal that is associated much more with irrelevant documents.

Table 6: Results of P@{10, 20}, nDCG and MAP

IR Models	P@10	P@20	nDCG	MAP
Baseline: Textual Model				
<i>LM.Hiemstra</i>	0.3700	0.3403	0.4325	0.2402
Baseline: Facebook Like				
<i>Like</i>	0.3938	0.3620	0.5130	0.2832
Facebook Reactions Individually				
<i>Love</i>	0.4122*	0.3702*	0.5300*	0.2978*
<i>Haha</i>	0.3900	0.3624	0.5100	0.2766
<i>Wow</i>	0.4031*	0.3755*	0.5203*	0.2889*
<i>Sad</i>	0.3800	0.3505	0.4811	0.2700
<i>Angry</i>	0.3111	0.2814	0.3421	0.1601
Facebook Reactions Combinations				
(<i>Love, Haha</i>) ⁺	0.4187*	0.3801*	0.5555*	0.2991*
(<i>Sad, Angry</i>) ⁻	0.3021	0.2614	0.3167	0.1574
(<i>Love, Haha, Wow</i>) ⁺	0.4275*	0.4112*	0.5773*	0.3168*

b) Impact of Grouped Reactions. The prior based on grouped reactions improve significantly the results in terms of nDCG compared to the *LM.Hiemstra*, especially when using positive reactions (*Love, Haha*)⁺: +29% and (*Love, Haha, Wow*)⁺: +34%, as well as compared to the consideration of reactions individually. We can also notice that when negative reactions are grouped, the relevance of the returned documents becomes very low. This amounts to a lack of topics that expresses appropriate emotions to *Sad* and *Angry*. In our case, we have only this two IMDb topics “Worst actor century” and “Chernobyl” that have registered a significant P@10 (0.7126 and 0.7610, respectively). This type of features can be helpful to capture non-relevant documents when the topics are positives. In addition, the reaction *Wow* is even more effective when considered with positive reactions group (*Love, Haha, Wow*)⁺: +34%, +10% against textual model and *like*-based model, respectively.

c) Correlation Between Reactions and Relevance. In order to analyze Facebook reactions and determine if there is a link (dependence/independence) between them and the document relevance, we conducted a correlation study using Spearman’s Rho correlation coefficient [4].

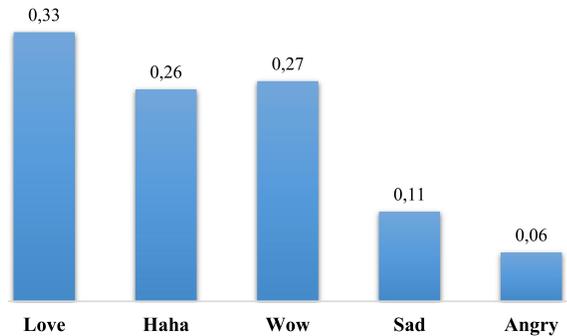


Figure 3: Rho Correlations of Facebook reactions

Figure 3 shows the values of correlations between ranges reactions with respect to documents relevance. This study shows that

Love (0.33) has the highest correlation, followed by number of *Wow* (0.27) and *Haha* (0.26). The reactions *Sad* and *Angry* are the less correlated with relevance. This results justifies the results obtained above (see table 6) and confirms that these novel emotional signals contribute to the improvement of the retrieval performance. Indeed, the well positioned resources have a high number of reactions and the frequently *loved* or *funniest* content is increasingly correlated with good ranking of relevance.

Finally, these results show that Facebook reactions are fruitful for IR systems, and can be more effective for specific search that take into account emotional aspect in the retrieval. Consequently, grouping these signals according to their meaning, where some signals are positive and related to the document reputation, is more effective compared to the individual consideration of signals to improve IR. We also note that the impact of these reactions is related to the nature and the emotion expressed in the topics.

6 CONCLUSION

This paper studied the impact of novel signals related to Facebook users’ reactions (*love, haha, angry, wow, and sad*) on IR. By analyzing these reactions, we notice that they present a positive growing. They allow to better understand if the user enjoyed the content or not. We proposed to estimate document priors by considering these reactions as an additional source of evidence to measure document relevance. Experimental evaluation conducted on IMDb dataset shows that taking into account these social features in a textual model improves the quality of returned search results. The correlation analysis shows that positive reactions are positively correlated with relevance.

For future work, we plan to estimate the impact of these reactions in sentiment detection with respect of their creation dates. Further experiments on another dataset are also needed. Unfortunately, until now, these social features are not yet available on other documents of standard dataset such as INEX Social Book Search.

REFERENCES

- [1] I. Badache and M. Boughanem. Harnessing Social Signals to Enhance a Search. In *ACM/IEEE WIC*, pages 303–309, 2014.
- [2] I. Badache and M. Boughanem. Social priors to estimate relevance of a resource. In *SIGIR IIX*, pages 106–114, 2014.
- [3] I. Badache and M. Boughanem. Fresh and diverse social signals: Any impacts on search? In *SIGIR CHIIR*, pages 155–164, 2017.
- [4] S. D. Bolboaca and L. Jantschi. Pearson versus spearman, kendall’s tau correlation analysis on structure-activity relationships of biologic active compounds. *Leonardo Journal of Sciences*, 2006.
- [5] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. In *WWW*, pages 107–117, 1998.
- [6] S. Chelaru, C. Orellana-Rodriguez, and I. Altingovde. How useful is social feedback for learning to rank youtube videos? *World Wide Web*, 2013.
- [7] BD. Davison. Topical locality in the web. In *SIGIR*, pages 272–279, 2000.
- [8] D. Hiemstra. A linguistically motivated probabilistic model of information retrieval. In *ECDL*, pages 569–584, 1998.
- [9] Y. Inagaki, N. Sadagopan, G. Dupret, A. Dong, C. Liao, Y. Chang, and Z. Zheng. Session based click features for recency ranking. In *AAAI Press*, 2010.
- [10] W. Kraaij, T. Westerveld, and D. Hiemstra. The importance of prior probabilities for entry page search. In *SIGIR*, pages 27–34, 2002.
- [11] C. Macdonald and I. Ounis. Voting for candidates: adapting data fusion techniques for an expert search task. In *CIKM*, pages 387–396, 2006.
- [12] C. Orellana-Rodriguez, W. Nejd, E. Diaz-Aviles, and IS. Altingovde. Learning to rank for joy. In *WWW, WWW ’14 Companion*, pages 569–570, 2014.
- [13] T. Westerveld, W. Kraaij, and D. Hiemstra. Retrieving web pages using content, links, urls and anchors. 2002.
- [14] BW. Zhang, XC. Yin, and F. Zhou. A generic pseudo relevance feedback framework with heterogeneous social information. *Information Sciences*, 2016.