

Diversifying User Comments on News Articles

Giorgos Giannopoulos^{1,2,*}, Ingmar Weber³, Alejandro Jaimes³, and Timos Sellis^{1,2}

¹ School of ECE, NTU Athens

² IMIS Institute, "Athena" Research Center

³ Yahoo! Research, Barcelona

Abstract. In this paper we present an approach for diversifying user comments on news articles. In our proposed framework, we analyse user comments w.r.t. four different criteria in order to extract the respective diversification dimensions in the form of feature vectors. These criteria involve content similarity, sentiment expressed within comments, article's named entities also found within comments and commenting behavior of the respective users. Then, we apply diversification on comments, utilizing the extracted features vectors. The outcome of this process is a subset of the initial comments that contains heterogeneous comments, representing different aspects of the news article, different sentiments expressed, as well as different user categories, w.r.t. their commenting behavior. We perform a preliminary qualitative analysis showing that the diversity criteria we introduce result in distinctively diverse subsets of comments, as opposed to a baseline of diversifying comments only w.r.t. to their content (textual similarity). We also present a prototype system that implements our diversification framework on news articles comments.

1 Introduction

The last years the concept of social web is growing exponentially. More and more users socialize through facebook, discuss current topics in forums, express their opinions/sentiments through blogs or twitter. The social web has also infiltrated in more traditional aspects of the web, such as news sites. Large corporations, like Yahoo! News¹, allow their users to comment on news articles, facilitating the aggregation and public exposure of a wealth of user contributed information and opinions. Although this feature itself contributes largely to the spread of information and promotes the freedom of expression, data management issues come up due to the large amount of information to be handled.

In our scenario, some news articles can gather tens of thousands of comments, which makes it impossible for interested users to review all of them. However, sometimes, the article's content itself is not enough for a user to form a complete view over a topic. The public opinion is a valuable resource that complements the article and represents

* This research has been co-financed by the European Union (European Social Fund - ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: Heracleitus II. Investing in knowledge society through the European Social Fund.

¹ <http://news.yahoo.com/>

the “wisdom of the crowds”. In this case, the user needs to be able to review a very small amount of as heterogeneous as possible comments, that represent different aspects of the article. On top of that, the user needs to be able, by selecting a specific comment, to review “similar” comments. Another use case scenario regards an archivist that needs to archive web information/resources about a specific topic. In this case too, the archivist should be able to “attach” to the primary resource (news article) complementary information (diverse user comments). This process would help, e.g., a future journalist that is trying to review past events, to gather as much diverse information on a topic as possible.

In this paper, we propose a set of criteria on which we adapt state of the art diversification algorithms in order to obtain a diverse subset of user comments on news articles. To the best of our knowledge, this is the first work handling the specific problem. The rest of the literature involves analysing user comments from several aspects, such as volume, political opinion, etc (see Section 6). On the other hand, diversification is, in most works, handled from the aspect of (mostly unstructured) search results.

The diversification criteria we propose consider the following comment features:

- Textual similarity. This is the baseline diversity criterion that is also used in the rest of the literature to diversify search results. The objective is to obtain comments with diverse content.
- Sentiment. We consider the sentiment of users expressed in the respective comments, w.r.t. the news article content. Sentiment is measured in a 9 grade scale $([-4, 4])$, expressing negative, neutral, or positive sentiment. The objective is to obtain comments covering the whole range of sentiments.
- Named Entities (NEs). We consider the Named Entities (Persons, Organizations, Locations) found in the news article. Then, for each comment, we examine which of these NEs are referred in its content. Again, the objective is for the selected set of comments to contain as many article’s NEs as possible.
- User Co-commenting behavior. We consider, for each user, all the articles she has commented on. Since each comment corresponds to one user, this information applies for comments too. Then, the objective is to select comments, so that their respective users have commented on as many different articles as possible.

After the above features are calculated for each comment, we apply an iterative algorithm that, at each step, compares the current, diversified set of comments with **all** the candidate ones. This comparison gives, for each diversity criterion a separate score. These scores are weighted and integrated into a final diversity score for each comment. However, diversity is not the only objective: although the problem requires that heterogeneous comments are gathered, these comments ought to be relevant to the initial news article. So, the final score of each candidate comment, at each step, is a weighted sum of its relevance score to the news article and its diversity score (distance) to the already selected comments.

We have conducted a preliminary qualitative analysis that demonstrates the effectiveness of our method, as opposed to diversifying comments only on content (textual similarity). We present an intuitive example that shows the importance of adding the previously mentioned criteria into the diversification process. The next step is to conduct

a thorough user evaluation study, where users will be asked to evaluate the compared comment sets in terms of novelty, interestingness and topic coverage, concepts closely related to the concept of diversity.

The remaining paper is organized as follows. In Section 2, we discuss some background information on diversification algorithms. In Section 3, we present our method for diversifying news articles comments. In Section 4 we shortly describe the implemented system. In Section 5, we present an example from our preliminary analysis, that demonstrates the effectiveness of the proposed method. Section 6 presents related work and, finally, Section 7 concludes and discusses further work.

2 Background

As analysed in Section 6, there are several works describing diversification algorithms. The main idea, however, for most of them, revolves around the same process. In what follows, we first describe two variations of this process, based on the analysis on dispersion objectives performed in [16] and, then, we demonstrate how we interpret these objectives in our diversification framework.

2.1 Baseline Objectives

The authors in [16] propose 3 diversification objectives which are closely related to the problems of maximum dispersion and facility dispersion [20], [21]. We focus on two of them, *MAXSUM* and *MAXMIN*, that can give us a better intuition about their difference in the diversification optimization process.

The first objective aims at maximizing the sum of the relevance and dissimilarity of the selected set. That is, maximizing the following function $f(S)$:

$$f(S) = (k - 1) \sum_{u \in S} w(u) + 2\lambda \sum_{u, v \in S} d(u, v) \quad (1)$$

where S is the set diverse items, $|S| = k$ is the number of diverse items (comments) required, $w(u)$ is the similarity score of item (comment) u to the respective resource (news article), $d(u, v)$ is the diversity score (distance) between items u and v and $\lambda > 0$ is a parameter specifying the trade-off between relevance and similarity.

The second objective aims at maximizing the minimum relevance and dissimilarity of the selected set. That is, maximizing the following function $f(S)$:

$$f(S) = \min_{u \in S} w(u) + \lambda \min_{u, v \in S} d(u, v) \quad (2)$$

2.2 Objectives Interpretation

Consider a set T of candidate comments to be selected into the k -size set S of diverse comments. The candidates are depicted as non-shaded circles in Figure 1. The shaded circles represent already selected comments that constitute the *current* diverse comment

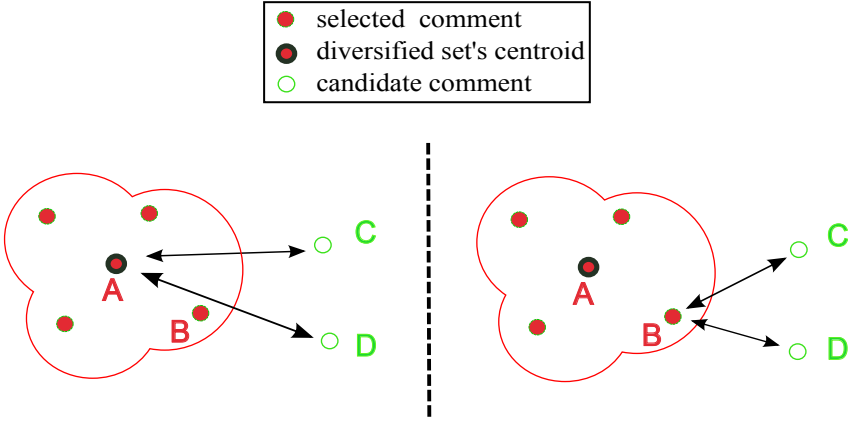


Fig. 1. Diversification Objectives

set. Both objectives described are based on the same iterative process: given the current set of diversified comments, for each candidate comment, produce a diversity score based on a certain score function. Then, select the comment that maximizes this score to insert into the diverse comments set.

The first objective (left figure) consists in selecting the candidate comment that has the maximum diversity distance to the current diverse set’s centroid. This way, in every iteration, each candidate is compared to the “average” of the current diverse set. The second objective (right figure) aims at finding the candidate comment that has the maximum distance from the closest comment belonging to the diversified set.

In our example, in the first case, candidates C and D are compared to the centroid of the diversified set, A , in terms of their diversity distance to it. This comparison results to D being the most distant candidate and, thus, the next comment to be inserted into the diverse set. In the second case, candidates C and D are compared to the closest comment of the diversified set, B , in terms of their diversity distance to it. This time, C is the most distant candidate and, thus, the next comment to be inserted into the diverse set. This example demonstrates that, different diversification methods may produce different solutions, that is, different sets of diverse comments. The proper diversification algorithm should be selected by considering the specific problem requirements.

3 News Comments Diversification

In what follows, we describe our method in detail. First, we introduce and analyse the diversification criteria we propose and, then, we describe how they are incorporated into the diversification process.

3.1 Diversification Criteria

As stated in Related Work Section, most works on diversification measure diversity in terms of content, that is textual (dis)similarity between items. Even in works where

more complex items are handled, e.g. [18] where items to be diversified are records with attributes, again, the distinct diversification criteria are defined on the textual similarity or matching of distinct attribute values. In this section, we extend the notion of diversity on new dimensions (apart from content) that include sentiment, named entities and user co-commenting.

Sentiment. We consider the sentiment expressed by users through their respective comments. We propose that sentiment (positive, negative, neutral) is a diversification factor, since it expresses users' opinions on the news articles' topics. In this sense, obtaining a set of comments that covers different classes of sentiment and, preferably, in a uniform manner, favors diversity.

We define 9 classes of sentiment within the interval $[-4, 4]$, with -4 denoting very negative sentiment, 4 very positive sentiment and 0 neutral sentiment. Each comment is assigned two different characterizations w.r.t. to the sentiment expressed within it:

- **Maximum/minimum sentiment.** We consider the whole text of the comment. Out of this, we extract the maximum positive sentiment, as well as the minimum negative sentiment value.
- **Average sentiment.** We regard each sentence of the comment separately and extract the respective positive and negative sentiment values. Then, we take the mean average of these values for all the sentences of the comment.

The sentiment extraction process is based on specific words found in the comment's text that express positive/negative sentiment. The above distinction into two types of extracted sentiment is performed in order to capture different facets of the expressed sentiment/opinion. For example, a comment may contain only one sentence that includes a very positive sentiment regarding a specific subtopic (e.g. a specific person) mentioned in the news article. On the other hand, the rest of the comment might be, on the whole, negative towards all the other aspects of the article. With the distinction we propose, we are able to capture these differences in sentiment expression.

After the sentiment extraction process, for each comment, we construct two 9-feature sentiment vectors, one for each type of sentiment extraction, with each feature corresponding to a different sentiment class. Each feature takes a boolean value that denotes whether the specific sentiment class is expressed in the comment.

Named Entities. We consider the Named Entities (NEs) found in the news article's text. These NEs might be Persons, Organizations or Locations. We suggest that NEs are important in terms of diversity, since news articles, most of the times revolve around NEs. Even when an article talks about events or situations, usually one or more Persons or Locations are involved. Given that, it is important for a diversified comment set to cover as many article's NEs as possible.

For each of the aforementioned NE categories, we create distinct NEs vectors, with each vector's features corresponding to the NEs found in the news article. For each comment, its feature values correspond to the frequency of the respective NE within the comment's text. In addition, we consider an aggregative NE vector that contains

all NEs, irrespective of category. This results to 4 NEs vectors, that represent, for each comment, the coverage of article’s Names Entities.

User Co-commenting Behavior. We consider the whole news article corpus. For each user, we construct a commenting vector, where each feature corresponds to a distinct news article. We suggest that the fact that a user comments a news article implies a relation between the user interests/opinions to the article’s topics. Also, the more comments the user has posed to an article, the more closely related to the article’s topic she is expected to be. Given that, the objective is to gather a diversified comment set, that corresponds to heterogeneous users. This heterogeneity, in our setting, is measured by the coverage of articles commented by the respective users.

For each comment, each feature value assigned is the commenting frequency of the user for the corresponding article, that is, the number of comments the user has posed on the article. For each user/comment, these feature values are normalized by the total number of comments the user has posed in **all** articles.

Content. Finally, we consider comments’ content, which is the baseline diversification criterion, used in most works handling diversification, e.g. in web search results diversification. The importance of comments’ content in the diversification process is straightforward. For each comment, we construct its term vector, with each feature corresponding to each distinct term found in the whole articles/comments corpus. Each feature value is computed by normalizing the term’s frequency within the comment by the total number of terms the comment contains.

3.2 Diversification Process

The previously described process produces, for each comment, four distinct vector categories, that map the comment to four diversification dimensions: Content, Sentiment, Named Entities and User Co-commenting Behavior. We perform diversification on each of the above dimensions separately, producing, each time, a separate diversity (dissimilarity) score for each comment. Then, these scores are aggregated into a final dissimilarity score that is their weighted sum.

In order to produce a diversity score, we need to define a diversity function that measures the distance between two items. We adopt the widely used cosine similarity score and we define the diversity score of two items, u, v , w.r.t. a specific dimension i , as:

$$d_i(u, v) = 1 - \cos_i(u, v)$$

where $\cos(u, v)$ is normalized in the interval $[0, 1]$.

However, diversity is not the only objective: although the problem requires that heterogeneous comments are gathered, these comments ought to be relevant to the initial news article. So, the final score of each candidate comment, at each step, is a weighted sum of its relevance score to the news article and its diversity score (as described above) to the already selected comments. We define the relevance score of a comment u , w.r.t.

the corresponding news article A , applying the cosine similarity measure on the article's and the comment's term vectors:

$$r(u, A) = \cos(u, A)$$

Depending on the diversification process we follow (as described in Section 2.2) we define two formulas that give the final score for each candidate comment u to be inserted into the set of diverse comments S , w.r.t. a news article A :

$$score_{MAXSUM}(u, A) = (1 - w) \cdot r(u, A) + w \cdot \sum_{i=1}^4 \lambda_i \cdot d_i(u, C_i)$$

where C_i is the centroid of the current diverse set w.r.t. the diversification dimension i , $w \in [0, 1]$ is the weight of the total diversity score, as opposed to relevance score and $\lambda_i \in [0, 1]$ is the weight of each individual diversity score, with $\sum_{i=1}^4 \lambda_i = 1$.

$$score_{MAXMIN}(u, A) = (1 - w) \cdot r(u, A) + w \cdot \sum_{i=1}^4 \lambda_i \cdot d_i(u, \min v_i)$$

where $\min v_i$ is the comment from the current diverse set that has the minimum distance to the closest comment from the candidate ones.

Below, the *MAXSUM* diversification algorithm is described. The *MAXMIN* algorithm is omitted since it is straightforward to derive it from the above formulas.

Algorithm 1. Produce diverse set of comments with MAXSUM

Input: Set of candidate comments T , size of diverse set k

Output: Set of diverse comments S

$S = \emptyset$

Insert into S the comment $u = \operatorname{argmax} r(u, A)$

while $|S| < k$ **do**

 Foreach dimension i , compute C_i

 Find the candidate comment $u = \operatorname{argmax} score_{MAXSUM}(u, A)$

 Insert u into S

end while

4 System Description

We divide the diversification process in two stages: Offline and online. Offline phase includes downloading the raw news articles and comments data, preprocessing them to extract feature vectors for the diversification dimensions, as well as term vectors for the relevance comparisons and storing them into the system's database. Online phase includes running the diversification algorithms on the extracted feature vectors.

Our system was implemented in Java. We used MySQL to store the preprocessed data. The data used in the specific process were downloaded from NY Times, using

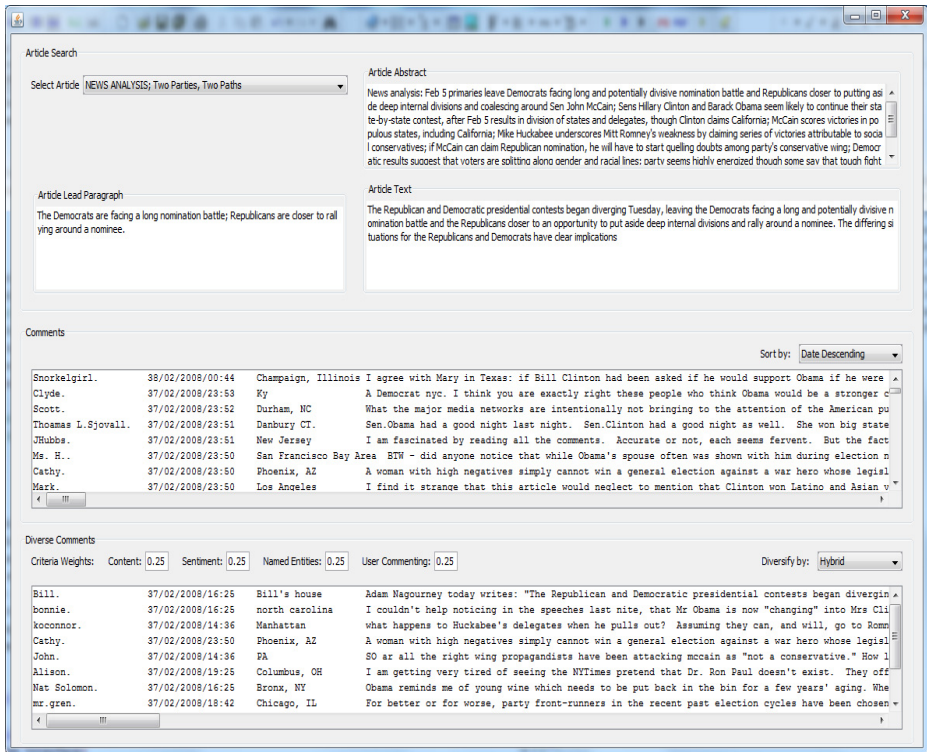


Fig. 2. System Interface

the respective APIs^{2, 3}. For sentiment extraction, we used SentiStrength⁴ ([23]) and for Named Entities recognition Stanford Named Entity Recognizer⁵ ([22]).

Figure 2 presents a screen of the implemented prototype. Through the upper panel ("Article Search") the user can select a news article and view the available information regarding it (text, abstract, lead paragraph). After an article is selected, its comments appear in the lower panels, sorted depending on the user choices. In the "Comments" panel, all article's comments are presented, either sorted by date, or sorted by their textual relevance to the article. In the "Diverse Comments" panel, a set of top-10 diversified results are presented according to the dimensions weight setup the user has chosen. In the specific example, all diversity dimensions are weighted equally. However, the user can select to diversify only by one dimension (e.g. sentiment) or even set the respective dimension weights on her own.

In the current prototype, we have set as default diversification algorithm the *MAXSUM* and default total diversity score weight $w = 0.5$. Of course, these choices

² http://developer.nytimes.com/docs/read/article_search_api

³ http://developer.nytimes.com/docs/community_api

⁴ <http://sentistrength.wlv.ac.uk/>

⁵ <http://nlp.stanford.edu/software/CRF-NER.shtml>

require further experimentation and tuning that will be part of our planned user evaluation study of the framework.

5 Evaluation

In this section, we present a preliminary analysis of the different diversification results achieved by our framework, as opposed to diversifying news articles comments only by their content. We note that this is a first cut study on the problem, that can only give us an intuition of the qualitative difference achieved. We are currently planning a thorough user evaluation study, where diversity of comment sets, produced by different approaches (ours, only content, only similarity to the article, etc), will be evaluated by users. In this setting, concepts like coverage, informativeness and novelty, that are closely related to diversity, will be evaluated.

In Table 2 (see last page) we present the abstract of a news article talking about the elections of US president candidates of two US parties, the 11-top diverse comments, diversifying only by content (baseline) and the 11-top diverse comments, diversifying by all the dimensions we propose. We chose 11 comments in each case, because, by default, we choose, for every method, the first comment to be selected, to be the most relevant by content to the news article. So, the first comment is always the same regardless the method used.

From a first look, we can see that the two methods have only 3 out of 10 comments in common. This proves, that, the extra dimensions we propose do matter in the diversification process. However, do they increase the quality of the resulting set, in terms of diversity? Defining the following facets on the specific article’s topic, we compare each set’s coverage on each of them: critique/analysis/opinion on (a)politicians, (b)media, (c)the election process, (d)voters. Table 1 gives us the coverage, in terms of distinct comments that handle the above issues, for each of the two compared sets. It is obvious that the comments set resulting by our method (Multiple criteria) covers in a more uniform way the several topic facets of the article. Also, most comments in the baseline diversified set involve criticizing or complaining about persons, parties or processes. While criticizing comments exist in our method’s comments set, there also exist comments where users perform analysis on the topic and try to give a better insight on the facts (see comments 2 and 8 in Multiple Criteria Diversification of Table 2). The above give us an intuition that our method covers, in general, more aspects of the article’s topic, as well selects comments of more heterogeneous objectives (critique, analysis, opinion).

Table 1. Topic facets and their coverage in comments(number of comments handling a facet) for baseline (Content) and our proposed method (Multiple criteria)

Facet	Content	Multiple criteria
Politicians	6	4
Media	1	4
Process	2	4
Voters	3	3

Table 2. Diverse comment sets comparison

Article's Abstract	
	Jared Diamond Op-Ed article holds that biggest global concern is resource consumption rate; notes that consumption rate in North America is 32 times higher than in developing world; says US promise that any country that adopts free-market economy can enjoy 'first-world lifestyle' is cruel hoax; contends that if China's per capita consumption rates rise to US levels, world will run out of resources at even faster rate; says it is futile to tell other countries not to reach for consumption rate that Americans already enjoy; contends that present rate of US consumption is unsustainable; says American consumption is wasteful and contributes little or nothing to quality of life; says that US consumption rates could be lowered if there was political will to tackle problem; drawing (M)
Content Diversification (Baseline)	
1	Adam Nagourney today writes: "The Republican and Democratic presidential contests began diverging Tuesday, leaving the Democrats facing a long...
2	I couldn't help noticing in the speeches last nite, that Mr Obama is now "changing" into Mrs Clinton. He implied he's now going to cover everyone with his healthcare policy program. That and virtually every other content point was Mrs Clinton's idea. What is he going to do, when Mrs Clinton is not around teaching him. Mr Obama (and Mr Kennedy and Ms Winfrey) – please step aside and let the people who know what they are doing, get to work. Talk is cheap and this presidency is not a job for the junior league. Thank you Bonnie Hauser
3	what happens to Huckabee's delegates when he pulls out? Assuming they can, and will, go to Romney, the latter may rightly see this still as a race.
4	A woman with high negatives simply cannot win a general election against a war hero whose legislative record is so indistinguishable from hers, particularly on the Iraq war.
5	SO ar all the right wing propagandists have been attacking McCain as "not a conservative." How long before they claim they never really meant it and start spewing lies for him and against the Democrats?
6	I am getting very tired of seeing the NY Times pretend that Dr. Ron Paul doesn't exist. They offer him almost no coverage – as they know that even bad press is good. The bias is so absolutely blatant; however, educated readers can read between the lines. How about some fair reporting. There are 4 Republican candidates, NY Times; report on all of them equally.
7	Obama reminds me of young wine which needs to be put back in the bin for a few years' aging. When it's more mature, it will probably be far more palatable and substantial.
8	For better or for worse, party front-runners in the recent past election cycles have been chosen more by default and less by way of solid qualification as leaders of vision. While I believe Tuesday's results are in general not all that surprising, the fact that the '08 cycle is unique and ground-breaking in many regards, is refreshing in and of itself. To what extent the much talked about "positive change" can come about when the dust finally settles and the rubber meets the road, of course, remains to be seen. That is the "Crux of the biscuit."
9	What the major media networks are intentionally not bringing to the attention of the American public is that Mike Huckabee's support is only partially derived from evangelical Christians, it's mostly derived from Mike's stated goal of ABOLISHING the INTERNAL REVENUE SERVICE, repealing the Sixteenth Amendment, and instituting the FairTax. This resonates with middle America as demonstrated in yesterday's voting. Many people I spoke with yesterday, who are NOT evangelical Christians by any stretch of the imagination, voted for Huckabee based ENTIRELY on his support of the FairTax, myself included. Get the facts: www.FairTax.org Then get the NYT #1 Bestselling Book by Neal Boortz and US Senator John Linder. WWW.FAIRTAX.ORG
10	I was surprised by the lack of astuteness on the voters' side. They are easily affected by youth and vacant idealism. A smart, effective Hillary Clinton should have been a LANDSLIDE. She has proven herself and her leadership, as well as her caring for all. I find it incredible that people want someone to be PRESIDENT WHO HAS NOT BEEN TRIED OR TESTED. NO WONDER OUR COUNTRY IS IN TROUBLE. IS APPOINTED IN RESULTS
11	John from Toronto: If Obama had attacked Bush and his policy with the same intensity and vigor with which he campaigns, then I would agree with you. Unfortunately, many Americans feel that he is a lot of bluster and inspiring phrases. These do not necessarily translate well to the Oval Office. Remember Cuomo's soaring speech? Cuomo? Ubi est?
Multiple Criteria Diversification	
1	Adam Nagourney today writes: "The Republican and Democratic presidential contests began diverging Tuesday, leaving the Democrats facing a long...
2	Many pundits are espousing the view that a long and drawn-out nomination battle between Obama and Clinton will turn negative, fracturing a Democratic Party that has found a unifying cause: anti-Bushism. However, there is a far more potentially divisive issue that looms ahead should the race go down to the wire: what is to be done with the polling results from Florida and Michigan? One can easily envision a scenario under which Obama obtains a razor thin majority of delegates heading into the nomination convention, but Clinton supporters claim she should be the party's nominee because Clinton has the greater number of delegates when the votes of Florida and Michigan are taken into account. It doesn't take much imagination to see a court case ensuing, proving once again that the Democratic Party is unequaled when it comes to snatching defeat from the jaws of victory.
3	When is the press going to start referring to the Democrats as being controlled by 'Left Wing Liberals' and 'non christians'. This would appear to be appropriate as the press seems to think the Republican party is composed of "Right Wing Conservatives" and "Evangelical Christians".
4	Why don't any of the columnists tell us how many democrats voted in the Tuesday primaries as opposed to republicans? This would give us an idea of the number of democrats and republicans who will be voting in the actual election.
5	Obama should withdraw because he doesn't have strength in the states we democrats need to win the november election. he is winning in red states. he is also winning in states with caucases which it is now clear are not reflective of the true voter sentiments and if he somehow where to get the nomination he would lead the party to defeat because his appeal is limited. it is time Ted Kennedy took him aside and explained reality to him. I think the voters in Mass wanted to send a reality check to Mr.Kennedy. some of us live in the real world and not in the magical kingdom inhabited by Obama,Oprah and the Kennedys
6	what happens to Huckabee's delegates when he pulls out? Assuming they can, and will, go to Romney, the latter may rightly see this still as a race. Excellent!! hope the media and the Obama fans take a good long breath and stop pushing Obama down our throats.Let us make up our minds,without the pundits incessant Hillary bashing.They are treating the primaries the way they treat Paris and Britney.Please grow up,guys and gal!
7	The results from AK, ID, UT, CO, KS, MO, ND and MN do not seem to tally back up claims by reporters such as your Mr. Adam Nagourney that a/the results suggest that Democrats are fracturing along gender and racial lines as they choose between a black man and a white woman.I am not aware of large African American populations in most of these states, yet Mr. Obama won them, and in most cases by as much as 60 percent of the vote. Why this continued obsession with gender and race when reporting about the Democrats? I am beginning to think reporters and pundits jump on these a'divisions because those are the easiest things to pick out that do not require serious analysis of voter's choices. Give us more insight into how people voted with respect to important issues such as the economy, healthcare, education, and foreign policy. Does it ever occur to you that people may actually select the candidate they vote for because of these issues and not because of their gender or race?
8	I am fascinated by reading all the comments. Accurate or not, each seems fervent. But the fact is that the Republicans are doing what they usually do: unifying; and the Dems are doing what they usually do: squabbling. I'm a Democrat, so I'm not criticizing (business as usual). Oh--and PS: I voted for Hillary. Anybody got a problem with that?
9	For better or for worse, party front-runners in the recent past election cycles have been chosen more by default and less by way of solid qualification as leaders of vision. While I believe Tuesday's results are in general not all that surprising, the fact that the '08 cycle is unique and ground-breaking in many regards, is refreshing in and of itself. To what extent the much talked about "positive change" can come about when the dust finally settles and the rubber meets the road, of course, remains to be seen. That is the "Crux of the biscuit."
10	Obama reminds me of young wine which needs to be put back in the bin for a few years' aging. When it's more mature, it will probably be far more palatable and substantial.
11	

6 Related Work

As stated at the Introduction, to the best of our knowledge, there are no works that can directly be compared with our proposed method. In what follows, we present several approaches that deal with the problems of (a) news comments analysis and (b) search results diversification.

The work in [11] is the closest to ours. The authors present ongoing work on a system regarding online discussion groups. The system first requires that users explicitly state their opinions of specific topics. Then, it exploits this feedback to recommend several opinions, allowing the user to vary the similarity/diversity degree of the recommendations, w.r.t. her own opinions. Apart from the difference in the diversity criteria used, the system described in [11] differs from ours in that it requires explicit, specific feedback from users and, also, it diversifies the recommended opinions w.r.t. each user's personal opinions and not in a global manner.

The authors of [7] propose a news recommendation system in forum-based social media, that exploits user comments to produce news recommendations. The approach aims at building a topic profile, utilizing both the news text and its comments. This profile is then used to retrieve relevant news articles. Similarly, [8] present a method for recommending to users news articles that are likely to be commented by them. The authors propose a hybrid recommendation approach, where they exploit, apart from document content, the co-commenting patterns of users on the respective articles.

In [2] the authors first predict whether a news article is to receive any comments at all and, then, whether it will receive many comments or not. To this end, they apply two separate classification phases. In [1] they try to model and compare commenting distributions from several news sources and, also, predict comment volume by observing a short first period of commenting.

In [4] the authors try to capture commenters sentiment patterns towards political news articles and to predict the political orientation from the sentiments expressed in the comments. The authors apply different learning techniques, depending on whether they predict political orientation for one or more commenters. They also take into account contextual information, such as the vote or links a comment received. In [12] the authors study user comments on political news and evaluate readers' satisfaction on political opinions. In this way, they aim to differentiate between users who seek similar opinions to theirs and users who seek diverse ones.

The authors in [6] study the descriptiveness of comments, i.e. the extent to which comments are similar to the topic they refer to. The authors obtain positive results, in the sense that a sufficient amount of comments can adequately represent the original commented text. In [3] the authors perform a study on users' needs w.r.t. news article comments and conduct a quality analysis on comments posed in the articles of an online newspaper. In [5] an analysis of links, comments and interconnections between blogs is performed. The authors of [9] aim at producing document summaries, utilizing the respective comments. To produce the summaries, they extract sentences from the original document (e.g. blog post), which are biased to keywords extracted from the document's comments. In [10] the authors perform an analysis on blog post comments and their relation to the posts. Specifically, they estimate the overall volume of comments in the blogosphere, analyze the relation between the weblog popularity and

commenting patterns in it and measure the contribution of comment content to weblog access.

A thorough review of fundamental works in diversification is given in [19]. [13] describe the maximal marginal relevance method, which attempts to maximize relevance while minimizing similarity to higher ranked documents. To this end, the relevance of search results is calculated using two similarity functions, one measuring the similarity among documents, and the other the similarity between document and query. [14] consider an evaluation metric that penalizes a retrieval model only if it retrieves no relevant results at all. Given that, they propose a method where each result document is selected based on the probability that it is relevant to the previously selected ones.

In [16], the authors introduce a set of diversification axioms and show that it is not possible for a diversification algorithm to satisfy all of them. Also, they propose three diversification objectives. These objectives differ in the level at which the diversity is calculated, e.g. whether it is calculated per separate document or on the average of the currently selected documents. The authors in [15] present a framework for evaluating novelty and diversity. Similarly, [17] propose a greedy diversification algorithm but, also, extend some state of the art IR evaluation measures, so that they can be used in the context of diversification. Finally, [18] present a method for efficient diversification of structured data, where the items to be diversified are not documents, but objects with distinct attributes (i.e. records in a database table).

7 Conclusion

In this paper, we presented a methodology for diversifying user comments on news articles. We introduced comment-specific diversification criteria and extended two state of the art diversification algorithms, so that they can incorporate these criteria. We implemented the above method into a prototype system that works on a publicly available news article dataset. A preliminary evaluation of our approach showed that it is, indeed, meaningful to go beyond textual similarity, when diversifying user comments on news article. Finally, the implemented framework is general enough to be adapted to other (similar) settings, such as forum discussions, tweets, etc.

Our future work lies on enhancing and extending the current diversification criteria, so that they can be integrated to each other and, thus, become more effective for the diversification process. For example, we consider combining the Sentiment and Named Entities criteria, so that we can extract and utilize sentiment values on specific Named Entities found in the news article and commented by users. Also, we plan to perform an extended user study, where users will be asked to evaluate the interestingness, novelty and topic coverage of the diverse comment sets produced by our method, as opposed to baseline comment sets (non-diversified, or diversified only on textual similarity).

Acknowledgements. This research is conducted as part of the EU project ARCOMEM⁶ FP7-ICT-270239.

⁶ <http://www.arcomem.eu/>

References

1. Tsagkias, M., Weerkamp, W., de Rijke, M.: News Comments: Exploring, Modeling, and Online Prediction. In: Gurrin, C., He, Y., Kazai, G., Kruschwitz, U., Little, S., Roelleke, T., R uger, S., van Rijsbergen, K. (eds.) ECIR 2010. LNCS, vol. 5993, pp. 191–203. Springer, Heidelberg (2010)
2. Tsagkias, E., Weerkamp, W., de Rijke, M.: Predicting the volume of comments on online news stories. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management (CIKM 2009), pp. 1765–1768 (2009)
3. Diakopoulos, N., Naaman, M.: Towards quality discourse in online news comments. In: Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW 2011), pp. 133–142 (2011)
4. Park, S., Ko, M., Kim, J., Liu, Y., Song, J.: The Politics of Comments: Predicting Political Orientation of News Stories with Commenters Sentiment Patterns. In: Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW 2011), pp. 113–122 (2011)
5. Herring, S.C., Kouper, I., Paolillo, J.C., Scheidt, L.A., Tyworth, M., Welsch, P., Wright, E., Ning, Y.: Conversations in the Blogosphere: An Analysis “From the Bottom Up”. In: Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS 2005), p. 107b (2005)
6. Potthast, M.: Measuring the descriptiveness of web comments. In: Proceedings of the 32nd International ACM SIGIR Conference on Research and Development (SIGIR 2009), pp. 724–725 (2009)
7. Li, Q., Wang, J., Chen, Y.P., Lin, Z.: User comments for news recommendation in forum-based social media. *Information Sciences: an International Journal* 180(24), 4929–4939 (2010)
8. Shmueli, E., Kagian, A., Koren, Y., Lempel, R.: Care to Comment? Recommendations for Commenting on News Stories. In: Proceedings of the 18th International Conference on World Wide Web, WWW 2012 (to appear, 2012)
9. Hu, M., Sun, A., Lim, E.: Comments-oriented document summarization: understanding documents with readers’ feedback. In: Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2008), pp. 291–298 (2008)
10. Mishne, G.A., Gance, N.: Leave a Reply: An Analysis of Weblog Comments. In: Proceedings of the WWW 2006 Workshop on Weblogging Ecosystem: Aggregation, Analysis and Dynamics, at WWW: The 15th International Conference on World Wide Web (2006)
11. Wong, D., Faridani, S., Bitton, E., Hartmann, B., Goldberg, K.: The diversity donut: enabling participant control over the diversity of recommended responses. In: Proceedings of the 2011 Annual Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA 2011), pp. 1471–1476 (2011)
12. Munson, S.A., Resnick, P.: Presenting diverse political opinions: how and how much. In: Proceedings of the 28th International Conference on Human Factors in Computing Systems (CHI 2010), pp. 1457–1466 (2010)
13. Carbonell, J., Goldstein, J.: The use of MMR, diversity-based reranking for reordering documents and producing summaries. In: Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 1998), pp. 335–336 (1998)
14. Chen, H., Karger, D.R.: Less is more: probabilistic models for retrieving fewer relevant documents. In: Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2006), pp. 429–436 (2006)

15. Clarke, C.L.A., Kolla, M., Cormack, G.V., Vechtomova, O., Ashkan, A., Büttcher, S., MacKinnon, I.: Novelty and diversity in information retrieval evaluation. In: Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2008), pp. 659–666 (2008)
16. Gollapudi, S., Sharma, A.: An axiomatic approach for result diversification. In: Proceedings of the 18th International Conference on World Wide Web (WWW 2009), pp. 381–390 (2009)
17. Agrawal, R., Gollapudi, S., Halverson, A., Jeong, S.: Diversifying search results. In: Proceedings of the Second International Conference on Web Search and Web Data Mining (WSDM 2009), pp. 5–14 (2009)
18. Vee, E., Srivastava, U., Shanmugasundaram, J., Bhat, P., Yahia, S.A.: Efficient Computation of Diverse Query Results. In: Proceedings of the 2008 IEEE 24th International Conference on Data Engineering (ICDE 2008), pp. 228–236 (2008)
19. Drosou, M., Pitoura, E.: Search result diversification. *ACM SIGMOD Record* 39(1), 41–47
20. Hassin, R., Rubinstein, S., Tamir, A.: Approximation algorithms for maximum dispersion. *Operations Research Letters* 21(3), 133–137 (1997)
21. Ravi, S., Rosenkrantz, D.J., Tayi, G.K.: Approximation Algorithms for Facility Dispersion. In: Gonzalez, T.F. (ed.) *Handbook of Approximation Algorithms and Metaheuristics*. Chapman & Hall/CRC (2007)
22. Finkel, J.R., Grenager, T., Manning, C.: Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling. In: Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL 2005), pp. 363–370 (2005)
23. Thelwall, M., Buckley, K., Paltoglou, G., Cai, D., Kappas, A.: Sentiment strength detection in short informal text. *Journal of the American Society for Information Science and Technology* 61(12), 2544–2558 (2010)