

Enhanced Information Access to Social Streams Through Word Clouds with Entity Grouping

Martin Leginus,¹ Leon Derczynski² and Peter Dolog¹

¹ Department of Computer Science, Aalborg University, Selma Lagerlofs Vej 300, 9200, Aalborg, Denmark

² Department of Computer Science, University of Sheffield, S1 4DP, Sheffield, United Kingdom

Keywords: Word Clouds, Recognized Named Entities, User Evaluation, Social Streams Access.

Abstract: Intuitive and effective access to large volumes of information is increasingly important. As social media explodes as a useful source of information, so are methods required to access these large volumes of user-generated content. Word clouds are an effective information access tool. However, those generated over social media data often depict redundant and mis-ranked entries. This limits the users' ability to browse and explore datasets. This paper proposes a method for improving word cloud generation over social streams. Named entity expressions in tweets are detected, disambiguated and aggregated into entity clusters. A word cloud is generated from terms that represent the most relevant entity clusters. We find that word clouds with grouped named entities attain significantly broader coverage and significantly decreased content duplication. Further, access to relevant entries in the collection is improved. An extrinsic crowdsourced user evaluation of generated word clouds was performed. Word clouds with grouped named entities are rated as significantly more relevant and more diverse with respect to the baseline. In addition, we found that word clouds with higher levels of Mean Average Precision (MAP) are more likely to be rated by users as being relevant to the concepts reflected. Critically, this supports MAP as a tool for predicting word cloud quality without requiring a human in the loop.

1 INTRODUCTION

A word cloud is a visual information access interface which presents prominent and interesting terms from the underlying data collection. Word clouds allow quick access and exploration over document collections (Kuo et al., 2007) and reduce information overload (Miotto et al., 2013). There are various studies about tag cloud generation from folksonomy data (Leginus et al., 2013; Venetis et al., 2011), but few studies available about word clouds generated from user generated content on social media (Leginus et al., 2015). To investigate information access over social media, we investigate the “model organism” of this data type, Twitter (Tufekci, 2014), a worldwide popular online social network where users publish daily an enormous amount of content (upwards of 600 million pieces of content per day). Therefore, Twitter users often face information overload while searching, browsing and exploring tweets (Bernstein et al., 2010). Improved information access through e.g. word clouds can reduce this information overload. For instance, the interactive browsing interface

Eddi, where a word cloud is a core component of the interface (Bernstein et al., 2010), helps to decrease information overload. According to users, Eddi gives a more efficient and enjoyable mode of browsing the enormous amount of user stream tweets. Similarly, word clouds ease e-health monitoring when browsing large collections of tweets (Lage et al., 2014).

Despite the benefits of word clouds for accessing and browsing social stream data, it remains a difficult type of text to handle. As a result of the diversity of language choice and spelling present in social media, end users are often presented with several different terms that refer to the same entity or concept, each term using different syntax and form. Hence, conventional methods of generating word clouds lead to undesirable results when applied to social stream text. In particular, variations of proper nouns create duplicated clusters, each of reduced prominence. Compounding the issue, variety in expression is increased by tight space constraints in some formats (like Twitter's 140-character limit) and by social media's generally informal, uncurated setting, as well as the inclusion of quasi-word hashtags (Derczynski et al., 2013;

Maynard and Greenwood, 2014). For example, the football club “Manchester United” may also be referred to as “MUFC” and “Man U”. Adding entries for each of these leads to a decrease in the prominence of this key concept, while also taking up space in the cloud and thus reducing its eventual diversity.

Redundancies in the word cloud might lead to user confusion and an inability to effectively browse, explore and retrieve other relevant content. Therefore, two aspects should be considered when designing word cloud generation algorithms. Primarily, one aims to condense divergent terms describing the same concept into a single term. Also, a cloud’s high-level diversity must be maximised, so the cloud gives a broad account of topics in the collection. These two conflicting requirements must be balanced to achieve an optimal word cloud.

The aim of this study is to improve word cloud generation by grouping co-referent entity expressions across multiple documents, applying existing named entity recognition systems in a novel fashion and grounding terms in this difficult genre to linked data resources. We systematically study the benefits of grouped named entities on word cloud generation. We use three established synthetic metrics – Coverage, Overlap and Mean Average Precision (MAP) – for word clouds generated from social media data (Venetis et al., 2011; Leginus et al., 2015). Further, to verify the findings of the synthetic evaluation, we perform a user study comparing clouds with grouped and un-grouped named entities. The main contributions and findings of this paper are:

- The best performing DivRankTermsEntities method significantly increases Coverage with respect to the baseline method ($p = 0.0363$) and significantly decreases Overlap ($p = 0.000094$) (decreased redundancies). In addition, access to relevant documents is improved.
- Word clouds with grouped named entities are significantly more relevant ($p = 0.00062$, one sample t-test) and diverse ($p = 0.003$, one sample t-test).
- Users report that word clouds with grouped named entities that attain higher levels of MAP are more relevant than the word clouds with decreased levels of MAP. Hence, the MAP metric should be considered and measured when designing word cloud generation methods.

The structure of the paper is as follows. Section 2 provides a brief description of relevant related work. Section 3 describes a general process of word cloud generation and points out the focus of our work, and presents a method for word cloud generation with grouped entities. In Section 4, we describe graph-based word cloud generation which is the underlying

framework for the later evaluation. Section 5 presents the findings from an offline evaluation of generated word clouds from TREC2011 microblog collection as well as results of the performed user study. Finally, we discuss the paper’s contributions as well as possible limitations of this work in Sections 6 and 7.

2 RELATED WORK

2.1 Word Cloud Generation

Tag cloud generation from folksonomy data has been thoroughly researched. Several tag cloud generation methods are proposed (Leginus et al., 2013; Venetis et al., 2011) and even synthetic metrics expressing tag cloud quality designed (Venetis et al., 2011). There are few studies that explore the benefits of word clouds for browsing social stream data. For instance, the browsing tool Eddi, where word cloud is a core component of the interface (Bernstein et al., 2010), helps to decrease information overwhelm. Similarly, word clouds are useful for the detection of epidemics when browsing thousands of tweets is needed (Lage et al., 2014).

Crowdsourcing has been used to recognize named entities in tweets (Finin et al., 2010). This study reports that word clouds with named entities recognized by human workers are considered better. This supports our motive to promote and improve the handling of named entities in word clouds recognized in tweets. Our contribution beyond the study from (Finin et al., 2010) is threefold. First, we perform grouping of recognized named entities, which has a positive impact on the generated word clouds. Second, we systematically study how to generate word clouds with named entities and measure performance using multiple metrics. Third, we compare these measured performances with user ratings, to discover relations between metrics and the user’s perspective.

2.2 Social Stream Text Reparation

Social stream text is noisy, and difficult to process with typical language processing tools (Derczynski et al., 2013). Consolidation of the varying expressions used to mention entities is possible, over large well-formed corpora (Hogan et al., 2012). Achieving this over social streams presents new challenges, in terms of the reduced context and heightened diversity of expression. We propose a simple consolidation technique and explore its positive impact on the word cloud generation. Other potential methods we could employ to improve cloud quality are normalization and coreference. Normalization (Han and Baldwin,

2011) applies to many low-frequency terms, and as a result has a low impact with named entities. Also, while normalisation can compare minor spelling mistakes, it typically does not condense highly orthographically different expressions of the same entity. Coreference requires context to operate – something that is absent in short social media stream messages. Mapping keywords to unambiguous entity references is difficult, but understood (Augenstein et al., 2013).

3 WORD CLOUDS GENERATION WITH GROUPED NAMED ENTITIES

The process of generating word clouds from social media data is comprised of several subsequent steps:

1. *Data collection* where underlying documents are aggregated with respect to a user query, profile or trending topic. Often, the whole document collection might be used for a word cloud generation. In this work, we aggregate tweets for word cloud generation with respect to a user query.
2. *Data preprocessing* where extracted terms or phrases can be clustered, lemmatized or normalized. Documents can be further enriched or annotated with recognized named entities. The aim of this work is to investigate how recognized named entities detected during this phase impact the following word cloud generation.
3. *Word cloud generation* where the most relevant and important terms from the underlying collection are selected and consequently a word cloud is generated. Different word selection methods can be applied (Venetis et al., 2011), (Leginus et al., 2013), (Leginus et al., 2015).

The goal is to explore how recognized and grouped named entities from the *Data preprocessing* phase affect consequent word cloud generation. Do grouped named entities improve the quality of word clouds in terms of Coverage, Overlap and enhanced access to relevant tweets? Which word cloud generation method gives best results when using named entities? We transform these research questions into the following two hypotheses:

- *H1: Word clouds with grouped recognized named entities improve Coverage, Overlap and Mean Average Precision of generated word clouds.*
- *H2: Word clouds with grouped recognized named entities are more relevant and more diverse with respect to a provided query from the user perspective.*

In the following, we describe a method for grouping recognized named entities from tweets.

3.1 Grouping Named Entities

Conventional named entity recognition is not sufficient due to the nature of Twitter data (Derczynski et al., 2015). Standard named entity recognition approaches do not perform well on tweets because of the error prone structure (misspellings, missing capitalization or grammar mistakes) and their short length. We propose a method that aims to recognize named entities, to link the possible aliases and consequently to generate a word cloud with the recognized and linked named entities. This method can be thought of as a *Data preprocessing* step when generating word clouds over data from social streams. We combine standard named entity recognition tools with linked data. Alternative names for recognized entities are exploited for term cluster creation for each named entity. A canonical term from an entity term cluster is selected and, if relevant and prominent enough, it is presented in the final word cloud. The method is summarized as follows:

1. Gather a tweet collection – a set of tweets corresponding to a certain trending topic or a query on Twitter.
2. Recognize named entities (**NER**) and disambiguate them (**entity linking**) using the TextRazor service, which performs this task relatively well (Derczynski et al., 2015).¹
3. Using linked data, find alternative names for the recognised entity. We used Freebase’s (Bollacker et al., 2008) *aliases* field for this. For instance, for the entity *Manchester United FC* the following aliases might be retrieved *Man United*, *Manchester United*, *Man Utd*, *MUFC*, *Red Devils*, *The Reds* or *United*.
4. Perform lemmatisation to group together all the inflicted forms of a word to exploit only the base form of the term.
5. Using the aliases, build a term cluster for each entity, containing e.g. *Manchester United*, *Man U*, *MUFC*.
6. Find canonical names, such as *Manchester United FC*.
7. Generate the “condensed” cloud with aggregated counts of entity mentioned frequencies with some word cloud generation technique.

This may be performed as a general-purpose technique, and also to “targeted” streams, e.g. where tweets are filtered based on user-defined criteria such as keywords or spatial regions.

¹See www.textrazor.com

4 GRAPH-BASED WORD CLOUD GENERATION

In this section, we describe a graph-based method for generating word clouds with and without entity grouping. The benefits of graph-based word cloud generation are following. First, the method identifies relevant and important keywords in underlying text collections. Several studies have empirically demonstrated the benefits of graph based methods over standard popularity or TF-IDF word cloud generation methods (Leginus et al., 2013; Leginus et al., 2015; Wu et al., 2010). Second, graph-based methods allow biasing of word cloud generation toward user preferences or search queries. Our graph-based selection methods firstly transforms terms space into a graph. Then, the stochastic ranking of vertices in the graph is performed. In this work, we consider only global ranking but the proposed methods can be easily applied to biased graph-based ranking e.g., biased towards a user query or a user profile.

4.1 Graph-based Creation

Extracted terms from underlying tweets are used to build a graph where each term is a graph vertex. If two terms (vertices) co-occur at least α times, we consider these two terms as similar. Eventually, for each similar term pair, two directed edges are generated $t_1 \rightarrow t_2$ and $t_2 \rightarrow t_1$. Hence, edges capture co-occurrence relations between individual terms.

4.2 Graph-based Ranking

Graph-based ranking of terms simulates a stochastic process i.e., random traversal of the terms in the graph. We use a PageRank-style algorithm (Leginus et al., 2013), but any other algorithm based on random traversal of the graph could be employed. The aim is to estimate the global importance of a term t . If needed, it is possible to bias ranking towards user preferences through a vector of prior probabilities \vec{p}_p . For global graph-based ranking, i.e. without introduced bias, we set each entry in $\vec{p}_p = \{p_1 \dots p_{|V|}\}$ to $\frac{1}{|V|}$ where V is the set of all graph vertices. The sum of prior probabilities in \vec{p}_p is 1. A random restart of stochastic traversal of the graph is assured with a back probability β which determines how often a random traversal restarts and jumps back to a randomly selected (following \vec{p}_p probability distribution) vertex in the graph. So, the β parameter allows adjustment of bias toward user preferences or to vertices that are globally relevant in the underlying graph. To simulate random traversal of the graph, iterative stationary

probability is defined as:

$$\pi(v)^{(i+1)} = (1 - \beta) \left(\sum_{u=1}^{d_{in}(v)} p(v|u) \pi^{(i)}(u) \right) + \beta \vec{p}_p \quad (1)$$

where $\pi(v)^{(i+1)}$ is a probability of visiting node v at time $i + 1$, $d_{in}(v)$ is the set of all incoming edges to node v and $p(v|u)$ is a transition probability of jumping from node u to node v . In this work, a transition probability is set to $p(v|u) = \frac{1}{d_{out}(u)}$ for nodes v that have an ingoing edge from node u , otherwise $p(v|u)$ equals 0. The resulting global rank of a term t after convergence is considered as relevance of t i.e.;

$$I(t) = \pi(t) \quad (2)$$

Top-k ranked terms are then used for word cloud generation where the ranking score indicates the prominence of the term in a word cloud.

5 EVALUATION

We retrieved available tweets with relevance judgements from TREC2011 microblog collection (Ounis et al., 2011) during August 2014. Although some tweets were not available during retrieval, we compare results over the same corpus. We do not consider the missing tweets as a limitation of our evaluation – see (McCreadie et al., 2012). The relevance judgements for TREC2011 microblog collection were built using a standard pooling technique. For TREC the relevance of a tweet with respect to a query was assessed with a three-point scale; 0: irrelevant tweet, 1: relevant tweet and 2: highly relevant tweet. In this work, we consider both relevant and highly relevant tweets as equally relevant.

5.1 Metrics

We evaluate individual aspects of generated word clouds using the synthetic metrics introduced in (Venetis et al., 2011; Leginus et al., 2015).

The generated word cloud with k terms is denoted as WC_k . A term t links to a set of tweets TW_t . TW_{t_q} is the set of all tweets that are associated with a query phrase t_q . The first metric is *Coverage*, defined as:

$$\text{Coverage}(WC_k) = \frac{|\cup_{t \in WC_k} TW_t|}{|TW_{t_q}|}, \quad (3)$$

where the numerator of the fraction is the size of the union set. The union set consists of tweets associated with each term t from the word cloud WC_k . $|TW_{t_q}|$ is the number of all tweets that are associated

with a query phrase t_q . The metric ranges between 0 and 1. When a Coverage for a particular word cloud WC_k is close to 1, the majority of tweets are “covered” i.e., linked from the word cloud WC_k .

Overlap of WC_k : Different words in WC_k may be linking to the same tweets. The Overlap metric captures the extent of such redundancy. Thus, given $t_i \in WC_k$ and $t_j \in WC_k$, we define the *Overlap*(WC_k) of WC_k as:

$$\text{Overlap}(WC_k) = \text{avg}_{t_i \neq t_j} \frac{|Tw_{t_i} \cap Tw_{t_j}|}{\min\{|Tw_{t_i}|, |Tw_{t_j}|\}}, \quad (4)$$

If *Overlap*(WC_k) is close to 0, then the intersections of tweets annotated by depicted words are small and such word clouds are more diverse.

Further, we measure Mean Average Precision metric (Leginus et al., 2015) for the evaluation of word clouds as follows:

1. For given terms and corresponding weights of a word cloud WC_k , create a query vector Q_{WC_k} with normalized weights. Each entry of the query vector Q_{WC_k} represents the importance of a term from the word cloud WC_k with the normalized weight i.e., more important terms from the word cloud are represented with higher weights.
2. Rank and retrieve top- k tweets matching a given query Q_{WC_k}
3. Measure mean average precision(MAP) where each relevant tweet from TREC2011 microblog collection is considered a positive.

Ranking of relevant tweets with respect to a given query Q_{WC_k} is computed with standard information retrieval function OKAPI BM25 which can be defined as:

$$S(tw, Q_{WC_k}) = \sum_{q_i \in Q_{WC_k} \cap tw} c(q_i, Q_{WC_k}) \cdot TF(q_i, tw) \cdot IDF(q_i) \quad (5)$$

where

$$TF(q_i, tw) = \frac{f(q_i, tw) \cdot (k_1 + 1)}{f(q_i, tw) + k_1 \cdot (1 - b + b \cdot \frac{|tw|}{\text{avg}twl})}$$

$$IDF(q_i) = \log \frac{N - n(q_i) + 0.5}{n(q_i) + 0.5}$$

and $f(q_i, tw)$ is a q_i term frequency within a tweet tw , $|tw|$ is the length of a given tweet tw , $\text{avg}twl$ is average length of tweet within the corpus, N is a total number of tweets in the corpus and $n(q_i)$ is the number of tweets that contain the term q_i . To capture the importance of a word from the generated word cloud, we multiply the whole relevance score for a given term with the word cloud weight $c(q_i, Q_{WC_k})$ for the given term q_i . The function $c(q_i, Q_{WC_k})$ returns

a weight of the term q_i from the query vector Q_{WC_k} which corresponds to the term weight from the word cloud WC_k . We set the same values for parameters $k_1 = 1.2$ and $b = 0.75$ as in (Manning et al., 2008).

We measured the average precision at K for the retrieved top K list of ranked tweets with respect to the given word cloud. Further, we measured the MAP for all generated word clouds. The average precision of top K ranked tweets with respect to the word cloud is calculated as follows:

$$AP@K(Q_{WC_k}) = \frac{\sum_k^K (P(k) \cdot \text{rel}(k))}{\#\text{relevanttweets}}$$

where $P(k)$ is the precision at k -th position in the ranked top K list and $\text{rel}(k)$ is 1 if the tweet at rank k is relevant, otherwise $\text{rel}(k)$ is 0 and $\#\text{relevanttweets}$ is the number of relevant tweets within the top K list. MAP is defined as:

$$MAP@K = \frac{\sum_{Q_{WC_k} \in AWC_k} AP@K(Q_{WC_k})}{|AWC_k|}$$

where AWC_k is the set of all generated word clouds and $AP@K Q_{WC_k}$ is average precision for the given word cloud Q_{WC_k} . In this work, we measure MAP at 30 under the assumption that it represents a reasonable cutoff for the number of relevant tweets similar to the approach in (Ounis et al., 2011).

5.2 Baseline Method

PageRank Exploiting Only Extracted Terms (PgRankTerms). This method was originally proposed in (Leginus et al., 2013) to estimate tag relevance wrt. a certain query, and it outperformed several tag selection approaches in terms of relevance. In this work, the method estimates global terms importance within the graph created from the pooled tweets for the individual query from TREC2011 microblog collection. The β parameter is set to 0.85 (recommended value for a Pagerank algorithm). Due to the short nature of tweets, threshold α for edge creation between individual terms is set to 0. Shorter texts lead to small numbers of co-occurring terms, which consequently leads to a sparse graph.

5.3 Entity Based Methods

Most Frequent Entities (MFE). This method selects only recognized entities as defined in Section 3.1. The method provides a list of entities sorted by frequency in descending order, selecting top- k most popular entities.

Most Frequent Entities with Grouped Aliases




Figure 1: Coverage, Overlap, and Mean Average Precision for word clouds of various sizes generated for queries from TREC2011 microblog collection.

(MFEA). This method selects only recognized entities and associated Freebase aliases as defined in Section 3.1. The method provides a list of entities sorted by frequency in descending order.

PageRank Exploiting Extracted Terms, Entities and Grouped Aliases (PgRankTermsEntities). This method estimates the global importance of terms and recognized named entities within the graph created from the extracted terms, recognized named entities and grouped Freebase aliases from pooled tweets for the individual query from TREC2011 microblog collection. The parameters are set to the same values as in the baseline method.

5.4 Results

We performed the evaluation on queries from TREC2011 microblog collection (Ounis et al., 2011). The MFE method has the worst Coverage ranging from 35% for word clouds with 10 terms to 45% for word clouds with 50 terms. MFEA has better Coverage with approximately 10% absolute improvement over the MFE method. The baseline method PgRankTerms attains greater Coverage than MFE and MFEA methods. The reason for higher Coverage of PgRankTerms is that entity mentions do not occur enough in tweets to outperform other extracted words. However, when extracted words are combined with grouped named entities like in PgRankTermsEntities, the improvements in Coverage are highest. The PgRankTermsEntities method outperforms all other word cloud generation methods. PgRankTermsEntities improves Coverage with respect to PgRankTerms and MFEA because it groups entity synonyms e.g. USA, US and America and represent them with the canonical entity name United States of America. In addition, it selects the most important terms which are not referring to named entities e.g., #service, #jobs for the query *BBC World Service staff cuts*. The relative improvements in comparison to PgRankTerms are 11% for 10 terms, 6% for 20 terms, 4% for 30

terms and 2% for 40 and 50 terms word clouds. Coverage improvements decrease as word clouds increase in size because the number of relevant/prominent recognized named entities in the underlying graph is lower. These results support the hypothesis H1: that grouping named entities improves the Coverage of word clouds.

Word cloud generation methods which exploit named recognized entities improve MAP. PgRankTermsEntities, MFE and MFEA outperform PgRankTerms in terms of MAP. The relative improvements of PgRankTermsEntities in comparison to PgRankTerms are 4% for 10 terms, 10% for 20 terms, 9% for 30 terms, 23% for 40 and 14% for 50 terms word clouds. Thus, word clouds with named recognized entities improve access to the relevant tweets of the corpus which validates the H1 hypothesis. The main reason for the attained improvements is that almost 89% of all relevant tweets from TREC2011 microblog collection contain at least one recognized entity. Similarly, 31% of all relevant tweets contain at least one Freebase alias (with minimal length of 4 characters). Comparing all pooled tweets from the TREC2011 microblog collection 77% contain recognized named entities and 28% of tweets contain at least one Freebase alias. Further, linking entity synonyms increases both Coverage and also the prominence of the named entity in the word cloud. Thus, it is more likely that the named entity will be represented in the word cloud and, if relevant for the query, it will improve access to the relevant tweets.

Improved access to relevant tweets and enhanced Coverage of word clouds can be attained through a combined selection of terms and recognized named entities. Thus, for enhanced word cloud generation it is important to combine recognized and grouped named entities with relevant and prominent terms from the underlying dataset.

The methods exploiting recognized named entities do have higher Overlap than the PgRankTerms method. We consider this finding interesting and

unanticipated. The increased redundancies in the generated word clouds are caused by imperfect NER tools. In particular, tweets with an ambiguous name entity such as *BBC News Service* link to several semantically similar entities such as *BBC*, *BBC News*, *BBC NEWS Service*, which might lead to higher Overlap scores. Further, detected Freebase aliases might often increase Overlap for the similar reason e.g., alias *us* for *United States* covers many irrelevant tweets. To minimize the impact of ambiguous aliases we restrict the alias detection to a minimum length of 4 characters and the alias may not be a stop word.

Lemmatisation also had a positive effect on word cloud generation. Lemmatising terms to group them improves Coverage 1.75% above the baseline, and 3% for the PgRankTermsEntities. Similarly, MAP improves with an increase of 11% for PgRankTermsEntities and 7% for the baseline technique. The negative impact of lemmatisation on word cloud generation is higher Overlap (decreased diversity of word clouds), with an increase of 3% using the baseline technique. As the result is overall positive, we included lemmatisation as a preprocessing step for all cloud generation methods.

5.5 Diversification

To overcome the problems introduced by higher redundancy in word clouds, we investigate how to maximize global relevance as well as diversity of selected terms. Instead of following greedy diversification approaches, we take a unified approach of ranking global relevance together with the diversification objective. We use the DivRank algorithm (Mei et al., 2010) which assumes that transition probabilities change over time following the “rich gets richer” principle. The transition probability from different vertices to a certain vertex is reinforced by the number of previous visits to that state. Hence, during a random walk, vertices with high weights are likely to consume the weights of their neighbors. Consequently, top ranked vertices tend to have low connectivity, which corresponds to more diversified ranking.

Figure 2 shows that with diversified word cloud generation, Overlap decreases. The relative improvements of DivRankTermsEntities outperforms the PgRankTerms baseline are 14% for 10 terms, 14% for 20 terms, 12% for 30 terms, 11% for 40 and 12% for 50 terms word clouds. The DivRankTermsEntities method significantly decreases Overlap in comparison to the PgRankTerms baseline (Wilcoxon signed-rank test, $p = 0.000094$) The improvements are even more significant with respect to PgRankTermsEntities method with 24% for 10




Figure 2: Overlap for diversified word clouds with the method DivRankTermEntities of various sizes generated for queries from TREC2011 microblog collection.

terms, 22% for 20 terms, 20% for 30 terms, 19% for 40 and 18% for 50 terms word clouds. In contrast, diversified word cloud generation significantly improves Coverage of word cloud generation. The improvement is statistically significant with respect to the baseline method PgRankTerms (Wilcoxon signed-rank test, $p = 0.0363$). The mean of relative improvements DivRankTermsEntities with respect to PgRankTermsEntities (the best performing method when measuring Coverage) is 2.35%.

Diversified word cloud generation from grouped and recognized named entities combined with extracted words decreases significantly Overlap, improves significantly Coverage and improves access to relevant tweets. This validates hypothesis *H1*.

5.6 Crowdsourced Evaluation

In order to verify the findings from empirical evaluation of word clouds with different synthetic metrics, we designed a crowdsourced user evaluation of generated word clouds. We selected 8 queries from TREC2011 microblog collection for which we generated word clouds with DivRankTermEntities and PgRankTerms methods. We included 4 queries where the enhancement of MAP for word clouds with named entities with respect to the baseline was the greatest (denoted as Impr. MAP). Similarly, we added 4 word clouds for queries where the Overlap has been decreased the most with respect to the baseline (denoted as Impr. diversity (\downarrow Overlap)). The answers sought by the user evaluation are twofold. First, are word clouds with named entities perceived as more relevant and diverse by the end users? Second, do measured synthetic metrics correlate with the ratings of relevance and diversity by users?

Participants were asked to view a pair of word clouds, a set of tweets related to a certain query, and




Figure 3: Green bins (ratings 4 and 5) in the histograms indicate positive rating towards word clouds with grouped named entities. Ratings 1 and 2 indicate user preference towards the baseline word clouds and rating 3 represents that the baseline and the word cloud with grouped entities are equally relevant or diverse.

a related Wikipedia article. Their task was to determine which word cloud was more relevant and which was more diverse. The user was asked to rate the relevance and diversity of an individual word cloud with respect to the query on a Likert scale of 1 to 5 (Rating 1: word cloud *A* is very relevant/diverse to the pertaining query; Rating 3 - both word clouds are equally relevant/diverse to the pertaining query; and Rating 5 - word cloud *B* is very relevant/diverse to the pertaining query). We altered assignment of word clouds with named entities to either word cloud *A* or *B* for each query to prevent user bias that “word cloud *A* (with named entities) is always more relevant and diverse”.

5.6.1 Non-grouped vs. Entity-grouped Clouds

Each word cloud pair was compared using 20 ratings from distinct users. For 7 out of 8 word clouds, the average ratings of relevance and diversity favoured word clouds generated with automatically grouped named entities. For simplicity’s sake, in the following we refer to word clouds generated with grouped entities as word cloud *B*; positive ratings are those over 3.0.

From 160 distinct relevance ratings, 89 were positive towards word clouds with named entities, 27 were neutral ratings and 44 were more towards the baseline generated word clouds (see Figure 3). Similarly for diversity ratings, 73 were positive towards word clouds with named entities, 51 were neutral ratings and 36 were more towards the baseline generated word clouds.

To further compare differences between word clouds generated by the baseline and clouds with grouped named entities, we performed a statistical significance test. The null hypothesis is that user ratings are normally distributed with mean 3.0, i.e., word clouds generated by the DivRankTermEntities and PgRankTerms methods are rated as equally relevant and equally diverse. For the relevance judgements, we found that word clouds generated by the DivRankTer-




Figure 4: Aggregated user ratings for three distinct groups of word clouds categorized according to the measured levels of synthetic metrics.

mEntities method are significantly better rated than the baseline word clouds ($p = 0.00062$, one sample t-test). Similarly, we determined that word clouds generated by the DivRankTermEntities are significantly better rated for diversity with respect to the baseline method ($p = 0.003$, one sample t-test). These findings support hypothesis *H2*: users find word clouds with grouped entities more relevant and diverse than those with no entity grouping.

Table 1: Three distinct groups statistics which were created according to the measured levels of synthetic metrics.

Group	# clouds	min δ	mean δ
Impr. MAP	4	0.14	0.26
Impr. diversity (↓ Overlap)	4	-0.02	-0.023
Decr. MAP & Overlap	2	-0.02	-0.133

5.6.2 Synthetic Metrics vs. User Perception

The second goal of the user evaluation is to determine whether word clouds with higher levels of measured synthetic metrics are rated by users as more relevant and diverse or vice versa. We focused on the MAP and Overlap metrics.² To determine the correlation

²Validation by users of the third metric introduced in (Venetis et al., 2011), Coverage, is only possible with an interactive user evaluation. Hence, we do not include “coverage assessment” of word clouds in this study.

between user judgements and synthetic metrics, we have created 3 different groups (see Table 1). We exploit the same two groups of word clouds Impr. MAP and Impr. diversity (\downarrow *Overlap*) as in Section 5.6.1. In addition, we added a group Decr. MAP & Overlap with two clouds where levels of MAP and Overlap were lower than the baseline word clouds. For each group, we report a minimum δ value which is a minimal difference between measured levels of the particular metric for word clouds with grouped entities and the baseline. Hence a minimum δ is a threshold of measured synthetic metric whether to include a word cloud into the particular group. For instance, the threshold $\delta = 0.14$ for the Impr. MAP group indicates that only those word cloud pairs where the improvements of MAP are at least 0.14 (comparing the baseline and DivRankTermEntities methods) are included. The mean of δ expresses the average value of differences in metric values for each word cloud pair in the group, e.g., the average improvements of MAP in the group Impr. MAP is 0.26. Note that negative values of δ reflect cases where the metric is lower than baseline. For Decr. MAP & Overlap group, we only report levels of MAP due to substantial differences in comparison to the Overlap levels which have very slight differences between the baseline and the word clouds with grouped entities.

When all the ratings aggregated altogether from three groups, word clouds with grouped entities are still rated significantly more relevant ($p = 0.0046$, one sample t-test) and diverse ($p = 0.00047$, one sample t-test) than the baseline.

The relation between created groups and user judgements is presented in Figure 4. Users rated word clouds with higher MAP as more relevant. Of 80 ratings, 46 (57.5%) indicated that word clouds with grouped named entities are more relevant than the baseline. Conversely, for the word clouds with the decreased MAP and Overlap, only 40% of the ratings indicate preference towards word clouds with grouped named entities. Hence, word clouds with higher MAP get 17.5% more positive ratings (4 or 5 ratings) than the baseline. The difference is even more pronounced for “rating 5 - much more relevant than the baseline word cloud”, where Decr. MAP & Overlap group attained only 7.5% from all ratings, whereas the Improved MAP group attained 18.75%. Therefore, we can conclude that word clouds with grouped named entities which attain higher levels of MAP are more likely to be better rated in terms of relevance by users.

When measuring diversity, word clouds from the Impr. diversity (\downarrow *Overlap*) group were slightly more rated as “equally or more diverse than word clouds generated by the baseline” than other groups. In par-

ticular, with Impr. diversity (\downarrow *Overlap*), we observed a decreased number of ratings, expressing that the baseline word cloud is much more diverse (3.75% for Impr. diversity (\downarrow *Overlap*) group and 12.5 for Impr. MAP). However, when looking at the decreased Map and Overlap group, the distribution of the ratings is fairly even. Hence, the Overlap metric is not a suitable predictor of user diversity ratings. This might be because the relative improvements of Overlap are too subtle to produce observable differences in user judgements of diversity.

On the other hand, 46.3% of word clouds with improved MAP and 45% of word clouds from Decr. MAP & Overlap group were rated as more diverse than the baseline. Therefore, users rating word clouds with grouped entities have tend to find them more diverse than word clouds with no grouping.

6 DISCUSSIONS AND LIMITATIONS

False positives during entity recognition may have reduced relevant ratings. For instance, a word cloud generated for the query “*Super Bowl, seats*” contained “*Super (2010 American film)*” which is irrelevant for this query. Similarly, for “*Kubica crash*”, the entity “*crash bandicoot*” ended up in the word cloud.

Some word clouds generated with the PgRank-TermsEntities suffered from increased Overlap. This was partially caused by imprecise named entity disambiguation where ambiguous named entities were not grounded correctly. Therefore, the quality of word clouds with grouped named entities is bounded by the precision of named entity annotation tools.

Evaluating word clouds with crowdsourced user evaluation is a challenging task due to uncertainty of reliability and quality of user ratings. In our pilot study, we aimed to ensure the quality of user ratings with pre-filtering quiz questions. However, we have observed that for test questions where users were asked to rate word cloud diversity (one cloud was supposed to be more diverse) many participants disagreed. Due to the subjective nature of the task, we disregarded a user “qualifying” phase (as is often best practice in crowdsourcing (Sabou et al., 2014)) and instead aimed to collect more user ratings and observe aggregated ratings. To further ensure the quality of the ratings, we accepted ratings only from participants in English-speaking countries, as word clouds were generated from tweets written in English.

7 CONCLUSIONS

Generating word clouds from social streams is a difficult task; users often discuss the same entity using multiple aliases. This leads to a direct degradation in the utility of word clouds for accessing this complex source of data. We proposed a technique that groups aliases of the same entity and represents them with a canonical term. The method improves the coverage of word clouds and access to the relevant content. Due to the imperfect nature of state-of-the-art named entity recognition methods, redundancy of terms in word clouds is often increased. Therefore, it is necessary to apply a method for diversifying terms. In this work, we found that the proposed technique not only significantly decreased redundancy but also attained significantly higher coverage than the baseline word cloud generation method, leading to better word clouds and therefore improved information access.

An extrinsic user evaluation supported our hypothesis that word clouds with grouped named entities are significantly more relevant and diverse than word clouds with no entity grouping. Further, word clouds with grouped named entities that attain higher levels of MAP are more likely to be rated as relevant by users.

Finally, it was shown that the previously-proposed MAP metric for automatic cloud evaluation predicts extrinsic human evaluations of cloud quality. Thus, when designing word clouds, the MAP metric should be used as a quality predictor of the cloud generation technique, enabling automatic assessment of word cloud quality without a human in the loop.

ACKNOWLEDGEMENTS

This work was partially supported by the European Union under grant agreement No. 611233 PHEME.

REFERENCES

- Augenstein, I., Gentile, A. L., Norton, B., Zhang, Z., and Ciravegna, F. (2013). Mapping keywords to linked data resources for automatic query expansion. In *Proceedings of the Second International Workshop on Knowledge Discovery and Data Mining Meets Linked Open Data*, pages 9–20.
- Bernstein, M. S., Suh, B., Hong, L., Chen, J., Kairam, S., and Chi, E. H. (2010). Eddi: interactive topic-based browsing of social status streams. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, pages 303–312. ACM.
- Bollacker, K., Evans, C., Paritosh, P., Sturge, T., and Taylor, J. (2008). Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of SIGMOD*, pages 1247–1250. ACM.
- Derczynski, L., Maynard, D., Aswani, N., and Bontcheva, K. (2013). Microblog-genre noise and impact on semantic annotation accuracy. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media*, pages 21–30. ACM.
- Derczynski, L., Maynard, D., Rizzo, G., van Erp, M., Gorrell, G., Troncy, R., Petrak, J., and Bontcheva, K. (2015). Analysis of named entity recognition and linking for tweets. *Information Processing & Management*, 51(2):32–49.
- Finin, T., Murnane, W., Karandikar, A., Keller, N., Martineau, J., and Dredze, M. (2010). Annotating named entities in twitter data with crowdsourcing. In *Proceedings of the Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, pages 80–88. ACL.
- Han, B. and Baldwin, T. (2011). Lexical normalisation of short text messages: Mkn sens a# twitter. In *Proceedings of ACL*, pages 368–378. ACL.
- Hogan, A., Zimmermann, A., Umbrich, J., Polleres, A., and Decker, S. (2012). Scalable and distributed methods for entity matching, consolidation and disambiguation over linked data corpora. *Web Semantics: Science, Services and Agents on the World Wide Web*, 10:76–110.
- Kuo, B. Y., Hentrich, T., Good, B. M., and Wilkinson, M. D. (2007). Tag clouds for summarizing web search results. In *Proceedings of WWW*, pages 1203–1204. ACM.
- Lage, R., Dolog, P., and Leginus, M. (2014). The role of adaptive elements in web-based surveillance system user interfaces. In Dimitrova, V., Kuflik, T., Chin, D., Ricci, F., Dolog, P., and Houben, G.-J., editors, *User Modeling, Adaptation, and Personalization*, volume 8538 of *Lecture Notes in Computer Science*, pages 350–362. Springer International Publishing.
- Leginus, M., Dolog, P., and Lage, R. (2013). Graph based techniques for tag cloud generation. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media*, pages 148–157. ACM.
- Leginus, M., Zhai, C., and Dolog, P. (2015). Personalized generation of word clouds from tweets. *Journal of the Association for Information Science and Technology*.
- Manning, C. D., Raghavan, P., and Schütze, H. (2008). *Introduction to information retrieval*, volume 1. Cambridge university press Cambridge.
- Maynard, D. and Greenwood, M. A. (2014). Who cares about sarcastic tweets? Investigating the impact of sarcasm on sentiment analysis. In *Proceedings of LREC 2014*, Reykjavik, Iceland. ELRA.
- McCreadie, R., Soboroff, I., Lin, J., Macdonald, C., Ounis, I., and McCullough, D. (2012). On building a reusable twitter corpus. In *Proceedings of SIGIR*, pages 1113–1114. ACM.
- Mei, Q., Guo, J., and Radev, D. (2010). Divrank: the inter-

- play of prestige and diversity in information networks. In *Proceedings of SIGKDD*, pages 1009–1018. ACM.
- Miotto, R., Jiang, S., and Weng, C. (2013). etacts: A method for dynamically filtering clinical trial search results. *Journal of Biomedical Informatics*, 46(6):1060–1067.
- Ounis, I., Macdonald, C., Lin, J., and Soboroff, I. (2011). Overview of the TREC-2011 microblog track. In *Proceedings of the 20th Text REtrieval Conference*.
- Sabou, M., Bontcheva, K., Derczynski, L., and Scharl, A. (2014). Corpus annotation through crowdsourcing: Towards best practice guidelines. In *Proceedings of LREC 2014*. ELRA.
- Tufekci, Z. (2014). Big questions for social media big data: Representativeness, validity and other methodological pitfalls. In *Proceedings of ICWSM*, pages 505–514. AAAI.
- Venetis, P., Koutrika, G., and Garcia-Molina, H. (2011). On the selection of tags for tag clouds. In *Proceedings of WSDM*, pages 835–844. ACM.
- Wu, W., Zhang, B., and Ostendorf, M. (2010). Automatic generation of personalized annotation tags for twitter users. In *Proceedings of ACL:HLT*, pages 689–692. ACL.

