

A Semantically Motivated Approach to Compute ROUGE Scores

Elaheh ShafieiBavani Mohammad Ebrahimi Raymond Wong

Fang Chen

University of New South Wales, Australia

Data61 CSIRO, Sydney, Australia

{elahehs, mohammade, wong, fang}@cse.unsw.edu.au

Technical Report
UNSW-CSE-TR-201706
February 2017

THE UNIVERSITY OF
NEW SOUTH WALES



School of Computer Science and Engineering
The University of New South Wales
Sydney 2052, Australia

Abstract

ROUGE is one of the first and most widely used evaluation metrics for text summarization. However, its assessment merely relies on surface similarities between peer and model summaries. Consequently, ROUGE is unable to fairly evaluate abstractive summaries including lexical variations and paraphrasing. Exploring the effectiveness of lexical resource-based models to address this issue, we adopt a graph-based algorithm into ROUGE to capture the semantic similarities between peer and model summaries. Our semantically motivated approach computes ROUGE scores based on both lexical and semantic similarities. Experiment results over TAC AEsOP datasets indicate that exploiting the lexico-semantic similarity of the words used in summaries would significantly help ROUGE correlate better with human judgments.

1 Introduction

Quantifying the quality of summaries is an important and necessary task in the field of automatic text summarization. Traditionally, this task involves a human assessment of various quality criteria (e.g. coherence, conciseness, grammaticality, informativity and readability) [11]. Therefore, manual evaluation requires a lot of time and expertise in the field of given texts. To tackle this issue, automatic evaluation metrics come into play. This advent opens a new door to meta-evaluation (i.e. evaluation of evaluation metrics [2]). On the importance of meta-evaluation and its impact on summarization research, Text Analysis Conference (TAC¹) provides the task of Automatically Evaluating Summaries of Peers (AESOP) to assess the correlation of evaluation metrics with human judgments.

Among the proposals for automatic evaluation metrics [8, 23, 5], ROUGE² [9] is the first and still most widely used one [6]. This metric measures the concordance of system-generated summaries (peer summaries) and human-generated reference summaries (model summaries) by determining n-grams, word sequences, and word pair matches. ROUGE has frequently been proven to correlate very well with human judgments [10, 16, 14]. However, its assessment heavily relies on surface similarities between peer and model summaries. Hence, it is unable to fairly evaluate abstractive summaries which might include semantically similar units with different lexical representations (e.g. paraphrasing).

For more clarity, consider the following two sentences: (i) *They strolled around the city*; (ii) *They took a walk to explore the town*. These sentences are semantically the same, but lexically different. If one of them is included in a model summary, while a peer summary contains another one, ROUGE or other surface based evaluation metrics cannot capture their similarity due to the minimal lexical overlap. Our aim is to help ROUGE with identifying the semantic similarities of linguistic items at the deepest sense level, and consequently tackling the main problem of its bias towards lexical similarities.

Considering senses instead of words, we make use of the Personalized PageRank (PPR) algorithm [7] to leverage repetitive random walks on WordNet 3.0 [3] as a semantic network, and obtain the probability distribution of each disambiguated sense over all senses in the network. The weights in this distribution denote the relevance of the corresponding senses. Our graph-based approach (GROUGE) favors semantic similarity scores between n-grams, along with their match counts (used originally in ROUGE), to perform both semantic and lexical comparisons of a peer summary text and a set of model summaries. We have made the package publicly available.

To demonstrate the effectiveness of our approach, we have conducted a set of experiments over the TAC 2010 and 2011 AESOP datasets. We have compared the output of GROUGE with three manual metrics of Pyramid, Readability, and Responsiveness. The results we have achieved via three metrics of correlation (i.e. Pearson, Spearman, Kendall) demonstrate that GROUGE variants significantly outperform their corresponding variants of ROUGE most of the time. Beyond just enhancing the evaluation prowess of ROUGE, this approach has the potential to expand the applicability of ROUGE to abstractive summarization as well. The rest of the paper is organized as follows. Section 2 summarizes the background. The proposed approach is explained in Section 3. Section 4 reports the utilized data, the performed meta-evaluation, and the achieved results.

¹<http://www.nist.gov/tac>

²Recall-Oriented Understudy for Gisting Evaluation

Finally, Section 5 concludes the paper.

2 Background

ROUGE includes a large number of distinct variants, including four methods of n-gram counting (ROUGE-N; S; W; L). In summarization literature, a few of these variants are often chosen arbitrarily to assess the quality of summarization approaches. ROUGE-1, ROUGE-2, and ROUGE-SU4 are reported to have a strong correlation with human assessments, and are frequently used to evaluate summaries [10, 16, 14]. ROUGE-1 and 2, respectively calculate unigram and bigram co-occurrence statistics. ROUGE-SU4 measures co-occurring bigrams with maximum skip distance 4. It is noteworthy that ROUGE-2 and SU4 have been defined as baseline systems in TAC AESOP task.

Although ROUGE is a popular evaluation metric, studies on improving the current evaluation metrics is still an open research area. Many of these efforts are analyzed and gathered in a survey provided by [22]. In this section, we try to briefly review the most significant ones. Since DUC 2005, the Pyramid metric [18] was introduced as one of the principal metrics for evaluating summaries in the TAC conference. However, this metric is semi-automated and requires manual identification of summary content units (SCUs). The approach proposed in [8] is based on comparison of basic syntactic units, so called Basic Elements (BE) between the peer and model summaries. This metric, namely BE-HM was specified as one of the baselines in the TAC AESOP task. Among participating systems in this task from 2009 to 2011, AutoSummENG [5] was reported as one of the top systems. This graph-based metric (DEMOKRITOSGR in the experiments), compares the graph representations of peer and model summaries.

Surface-based evaluation metrics work well as long as a surface-based summary (i.e. extractive) is to be assessed. Difficulties arise while evaluating abstractive summaries including terminology variations and paraphrasing. For example, consider the following two phrases [13]: (i) *It is raining heavily*; (ii) *It is pouring*. If we are performing a lexical string match, as ROUGE does, there is nothing in common between the terms "raining", "heavily", and "pouring". However, these two phrases are semantically the same. Hence, we study the effectiveness of semantically motivated approaches to measure word semantic similarity on improving ROUGE evaluation. For this purpose, approaches can be grouped into two categories of distributional and lexical resource-based [20]. A recent branch of distributional models uses neural networks to directly learn the expected context of a given word and model it as a continuous vector [24, 1], often referred to as word embedding. In the context of summarization evaluation, an automated variant of the Pyramid metric which uses word embeddings to map text content within peer summaries to SCUs has recently been proposed [17]. However, the SCUs still need to be manually identified. To overcome this deficiency, a more recent automatic metric [13], namely ROUGE-WE has enhanced ROUGE by incorporating the use of a variant of word embeddings, called word2vec [12]. However, a good performance for Word2vec is usually obtained upon tuning different configurations of this model on a large number of different datasets [1].

Lexical resource-based approaches usually make an assumption that the similarity of two words can be calculated in terms of the similarity of their closest senses. Among them, a random walk-based method that models disambiguated words through the distributions of the PPR algorithm on the WordNet graph has proven to be promising [20]. Unlike this approach, none of the above-mentioned techniques disambiguates the words being compared, and they hence consider a word as a conflation of all its

meanings, which potentially reduces the quality of similarity measurement. Therefore, we are prompted to disambiguate n-gram pairs to a set of intended senses prior to modeling. This will make us able to identify the semantic similarity of peer and model summaries, independently of their surface forms or any semantic ambiguity therein.

Given a pair of peer and model summaries, we first utilize the PPR algorithm to acquire probability distributions of their words’ senses over the WordNet graph (PPR vectors). Comparing these vectors obtained for all senses in a pair of peer and model summaries, we disambiguate each word into its most appropriate sense. This helps us to measure the semantic similarity of n-grams at the deepest sense level. PPR vector is calculated this time for each of the model-gram (each n-gram in the model summary) and the peer summary text by initializing random walks from their disambiguated senses over WordNet. We further compute their semantic similarity by comparing the resulting PPR vectors. This remedy is finally adopted into ROUGE variants, and the appropriate weights of lexical and semantic similarity scores are explored through our experiments.

3 The Proposed Approach

ROUGE assumes that a peer summary is of high quality if it shares many words or phrases with a model summary. However, different terminology may be used to refer to the same concepts and hence relying only on lexical overlaps may underrate content quality scores. To tackle this issue, our approach utilizes both semantic and lexical similarities between a peer and its corresponding model summary. This method also enables us to reward terms that are not lexically equivalent, but semantically related.

3.1 Measuring Semantic Similarities

Given a pair of peer and model summaries, we need to compute and compare PPR vectors at the following levels: (i) *sense level*, to disambiguate each word (having a set of senses); and (ii) *n-gram level*, to measure the semantic similarity. Next, we explain how a PPR vector is constructed for a sense or a set of senses, and how a similarity score is computed accordingly.

PPR Vectors: To construct a PPR vector, we perform iterative random walks beginning at a sense (seed) or a set of senses (seeds) on WordNet. This provides a frequency or multinomial distribution over all senses in WordNet [19]. A higher probability will be assigned to senses that are frequently visited from the seeds. This representation is applicable both when the item itself is a single sense and when the item is a sense-tagged text. For better clarity, consider an adjacency matrix W for the WordNet graph, where edges connect senses according to the relations defined in WordNet (e.g. hypernymy, synonymy, etc.). The probability distribution for the starting location of the random walker in the network is denoted by $\vec{V}^{(0)}$. Given the set of senses S in a lexical item, the probability mass of $\vec{V}^{(0)}$ is uniformly distributed across the senses $s_i \in S$, with the mass for all $s_i \notin S$ set to zero. The PPR vector is then computed using Equation 3.1.

$$\vec{V}^{(t)} = (1 - \alpha)W\vec{V}^{(t-1)} + \alpha\vec{V}^{(0)} \quad (3.1)$$

where at each iteration, the random walker may jump to any node $s_i \in S$ with probability $\alpha/|S|$. Following the standard convention, the value of α is set to 0.15. The

number of iterations is also set to 30, which is sufficient for the distribution to converge. The resulting probability vector $\vec{V}^{(t)}$ is the PPR vector of the lexical item, as it has aggregated its senses' similarities over the entire graph. The UKB¹ implementation of PPR is used to this end.

To compare the PPR vectors of each pair of n-grams, we use an effective method, namely Weighted Overlap, which has consistently proven to be superior to cosine similarity, Jensen-Shannon divergence, and Rank-Biased Overlap for comparing vectors in different datasets [20]. This algorithm first sorts the two vectors according to their values and then harmonically weights the overlaps between them. Finally, the semantic similarity (Sim_{sem}) of two vectors V_1 and V_2 is calculated by Equation 3.2.

$$Sim_{sem}(V_1, V_2) = \begin{cases} \frac{\sum_{h \in H} (r_h(V_1) + r_h(V_2))^{-1}}{\sum_{i=1}^{|H|} (2i)^{-1}}, & \text{if } |H| > 0 \\ 0, & \text{otherwise} \end{cases} \quad (3.2)$$

where H denotes the intersection of all senses with non-zero probability (dimension) in both vectors, and $r_h(V_j)$ denotes the rank of the dimension h in the sorted vector V_j , where rank 1 denotes the highest rank. The denominator is also used as a normalization factor that guarantees a maximum value of one. The minimum value is zero and occurs when there is no overlap between the two vectors, i.e. $|H| = 0$. Next, we explain the process of n-gram disambiguation into a set of appropriate senses.

Disambiguation of n-grams: Prior to measuring semantic similarities, each word in the n-grams has to be analyzed and disambiguated into its intended sense. However, conventional word sense disambiguations are not applicable due to the lack of contextual information. Hence, we make use of an alignment-based algorithm proposed by [19] to disambiguate each word. This algorithm seeks the semantic alignment that maximizes the similarity of the senses of the compared words. In our approach, given two n-grams, for each word type t_i in n-gram G_1 , this algorithm assigns t_i to the sense that has the maximal similarity score to any sense of the word types in the compared n-gram G_2 . As an example, let us consider two sentences of "a1. Officers fired." and "a2. Several policemen terminated in corruption probe.", the semantic alignment procedure has been performed as follows [20]:

$$\begin{aligned} P_{a1}. & \text{officer}_n^3, \text{fire}_v^4 \\ P_{a2}. & \text{policeman}_n^1, \text{terminate}_v^4, \text{corruption}_n^6, \text{probe}_n^1 \end{aligned}$$

where P_j denotes the corresponding set of senses of sentence j . t_p^i denotes the i -th sense of a word t in WordNet with part-of-speech p . After alignment, among all possible pairings of all senses of $fire_v$ to all senses of all words in $a2$, the sense $fire_v^4$ (employment termination) obtains the value $Sim_{sem}(fire_v^4, terminate_v^4) = 1$, which is the maximal similarity value. Sim_{sem} gives the semantic similarity of two senses by comparing their PPR vectors, as defined in Equation 3.2. Therefore, GROUGE transforms the task of determining overlapping n-grams in ROUGE into that of computing the similarity of the best-matching sense pair across the two n-grams. It also enables the same n-grams to have different meanings when paired with different linguistic items. In the following, the generated PPR vectors for a pair of disambiguated model-gram and peer summary are compared to calculate their semantic similarities.

¹<http://ixa2.si.ehu.es/ukb/>

Model-grams against Peer Summary: Exploiting underlying semantic similarities between all n-gram pairs in the model and peer summary texts takes a lot of time and effort. To overcome this issue, we consider the peer summary text as a sense-tagged unit, and measure its semantic similarity against each n-gram in the model summary text (model-gram). For better clarity, let us consider $MT = \{mt_1, mt_2, \dots, mt_n\}$, and $PT = \{pt_1, pt_2, \dots, pt_m\}$ as the sets of tokens of a model and a peer summary text, respectively. Figure 3.1 shows how PPR vectors of unigrams and bigrams in a model summary text are compared to the PPR vector of the peer summary text.




Figure 3.1: Comparing PPR vectors between model-grams (unigrams and bigrams in a model summary text) and a peer summary text

Measuring semantic similarities and sense disambiguation are previously explained in details. We can list the steps as follows: (i) Generating PPR vectors for all senses in the model-gram and peer summary text; (ii) Comparing the PPR vectors to disambiguate the model-gram and peer summary text to a set of proper senses; (iii) Generating one PPR vector for each of the model-gram and peer summary text by initializing random walks from their disambiguated senses over WordNet; (iv) Comparing the resulting PPR vectors to compute the semantic similarity between the model-gram and peer summary text. Treating the peer summary text as one unit not only reduces comparison time and increases the efficiency, but also provides a suitable number of content words which guarantees implicit word sense disambiguation, and semantic relationship derivation.

3.2 OOV Handling

Similarly to any other graph-based approach that maps words in a given textual item to their corresponding nodes in a semantic network, modeling n-grams through PPR vectors can suffer from the limited coverage of words. This means that only those words that are associated with some nodes in WordNet can be handled. Since out-of-vocabulary (OOV) words are the words that are not defined in the corresponding lexical resource, they will be ignored while generating PPR vectors. The reason is that they do not have an associated node in the WordNet graph for the random walk to be initialized from. Denying OOV words, such as infrequent named entities, acronyms or jargon, while increasing in a text, can be problematic when measuring semantic similarity of n-gram pairs. To take OOV words into consideration, we follow the approach proposed by [20] and directly insert each OOV word into the resulting PPR vector. To this end, we introduce new dimensions in the resulting PPR vector, one for each OOV term, while assigning a weight to the new dimension so as to guarantee its placement among the top dimensions in its PPR vector.

3.3 Multiple Levels of Evaluation

Most single automatic metrics use one level of evaluation (i.e. lexical, syntactic or semantic). A better approach is to assess the results while combining multiple levels of evaluation into one model [2]. For better clarity, consider the following groups of sentences:

- a1. Soldiers were killed.*
- a2. Soldiers were executed.*
- a3. Military personnel were executed for committed crimes.*

- b1. Soldiers were killed.*
- b2. Soldiers were murdered.*
- b3. Several servicemen were murdered by criminals.*

Surface-based approaches that are merely based on string similarity cannot capture the similarity between any of the above pairs of *a1* and *a3*, or *b1* and *b3* as there exists no lexical overlap. In addition, a surface-based semantic similarity approach considers both *a1* and *b1* as being identical sentences, whereas we know that different meanings of the verb "kill" are triggered in the two contexts. Although verbs "kill", "execute" and "murder" are close together in WordNet, *a2* and *b2* carry very different connotations. As a remedy, we need to transform words to senses and perform disambiguation by taking into account the context of the paired linguistic item, hence providing a deeper measure of similarity comparison. We finally combine the lexical and semantic similarity scores to calculate GROUGE-N (Equation 3.3). This approach can increase the chance of getting the evaluation results more correlated with human assessments.

$$\begin{aligned} \text{GROUGE-N} &= \frac{\sum_{S \in \{\text{ModelSummaries}\}} \sum_{gram_n \in S} Sim_{LS}(gram_n, P)}{\sum_{S \in \{\text{ModelSummaries}\}} \sum_{gram_n \in S} Count(gram_n)} \end{aligned} \quad (3.3)$$

where n stands for the length of n-gram, and Sim_{LS} is the score of lexico-semantic similarity between a model-gram ($gram_n$) and the peer summary text (P).

To compute Sim_{LS} , we have conducted a set of experiments using lexical similarities, $Count_{match}(gram_n, P)$, and/or semantic similarities, Sim_{sem} (Equation 3.2). Note that $Count_{match}(gram_n, P)$ is the maximum number of n-grams co-occurring in a peer summary and a set of model summaries. The best correlation is obtained while using a linear combination of both scores with different weights according to Equation 3.4.

$$\begin{aligned} Sim_{LS}(gram_n, P) &= \beta \times Count_{match}(gram_n, P) \\ &+ (1 - \beta) \times Sim_{sem}(gram_n, P) \end{aligned} \quad (3.4)$$

The scaling factor β was optimized on the TAC 2010 AESOP dataset [15], and set to 0.5 to reach the best correlation with the manual metrics of Pyramid and Responsiveness.

4 Experiments

4.1 Data and Meta-evaluation

For the task of summarization evaluation, TAC has provided two benchmark AESOP datasets (AESOP 2010 and 2011), on which we can assess GROUGE. We make use of the TAC 2010 AESOP dataset to optimize the scaling factor, and the TAC 2011 AESOP dataset to evaluate GROUGE. This dataset consists of 44 topics, and two sets of 10 documents for each topic: set A for initial summaries; set B for update summaries. There are four human-crafted model summaries for each document set. A summary for each topic is generated by each of the 51 summarizers which participated in the main TAC summarization task. Source documents for summarization are taken from the New York Times, the Associated Press, and the Xinhua News Agency newswire.

Two different types of evaluation were tasked in TAC 2011 AESOP: *All Peers* and *No Models*. The former case assigns a score to each peer summary, including the model summaries. This evaluation is intended to focus on whether an automatic metric can distinguish between human and automatic summarizers. The latter assigns a score to each peer summary, excluding the model summaries. This case is intended to focus on how well an automatic metric is able to assess automatic summaries. Using model summaries as references, each automatic summary can be evaluated against all four references simultaneously. Since our aim is to evaluate the quality of automatic summaries, we make use of *No Models* evaluation.

The output of participating automatic metrics is tasked to be compared against human judges using three manual metrics of *Pyramid*, *Readability*, and *Responsiveness*. Hence, the outputs are scored based on their summary content, linguistic quality, and a combination of both, respectively. Prior to computing correlation of GROUGE variants with manual metrics, GROUGE scores have reliably been computed (95% confidence intervals) under ROUGE bootstrap resampling with the default number of sampling point = 1000. Correlation of GROUGE evaluation scores with the human judgments is then assessed with three metrics of correlation: Pearson r ; Spearman ρ ; and Kendall τ .

The value of all measures is between -1 and 1 of which 1 or -1 indicates a strong relationship between the two measures. The closer the value is to zero, the weaker the relation between the two measures. 25 automatic metrics participated in the TAC 2011 AESOP task, three of which (i.e. ROUGE-2, ROUGE-SU4, and BE-HM) were used as baselines. In our experiments, the effectiveness of GROUGE is demonstrated by assessing its three variants (GROUGE-1, 2, and SU4) against their corresponding variants of ROUGE, and the other 23 AESOP participants. Note that ROUGE-1 was not among the participating metrics, but will be considered in our experiments. We compute scores using the default NIST settings for baselines in the TAC 2011 AESOP task (with stemming and keeping stopwords¹).

4.2 Results

We have conducted a set of experiments to evaluate three variants of GROUGE (i.e. GROUGE-1, 2, and SU4), against the top 13 best-performing metrics among the 22 metrics participated in AESOP, the baselines (i.e. ROUGE-2, SU4, BE-HM), ROUGE-1, and the most recent related work (ROUGE-WE). Correlation results of the best-performing AESOP metrics with Pyramid, Responsiveness, and Readability scores to the correlation metrics of Pearson r , Spearman ρ , and Kendall τ are depicted in Figures

¹https://tac.nist.gov/2011/Summarization/AESOP.2011_guidelines.html

4.1, 4.2, and 4.3, respectively. The highest correlation results are highlighted for better clarity. To demonstrate the effectiveness of GROUGE in the ROUGE framework, the obtained correlation results of all variants of ROUGE-based metrics (ROUGE, ROUGE-WE, and GROUGE) with Pyramid, Responsiveness, and Readability are provided in Tables 4.1, 4.2, and 4.3, respectively. The best correlation in each column has been specified in bold.

Analyzing the correlation results obtained by the best-performing AESOP metrics in Figure 4.1 show that GROUGE-2 achieves the best correlation with Pyramid, regarding the Spearman and Kendall rank correlations. However, ROUGE-SU4 displays the best correlation with Pyramid considering the Pearson correlation. The key difference between the Pearson correlation and Spearman/Kendall rank correlation, is that the former assumes that the variables being tested are normally distributed, and linearly related to each other. The latter two measures are however non-parametric and make no assumptions about the distribution of the variables being tested. The assumption made by the Pearson correlation has been known too constraining [13], given that any two independent evaluation systems may not exhibit linearity.




Figure 4.1: Correlation of the **best-performing AESOP metrics** with the manual metric of Pyramid using the correlation metrics of Pearson r , Spearman ρ , and Kendall τ on the TAC 2011 AESOP dataset

Looking closer to the correlation with Pyramid scores, obtained by the variants of ROUGE-based metrics in Table 4.1, we observe that every GROUGE variant outperforms its corresponding ROUGE and ROUGE-WE variants, regardless of the correlation metric used. However, the only exception is ROUGE-SU4, which correlates slightly better with Pyramid when measuring with Pearson correlation. One possible reason is that Pyramid measures content similarity between peer and model summaries, while the variants of GROUGE favor semantics behind the content for measuring similarities. Since some of the semantics attached to the skipped words are lost in the construction of skip-bigrams, ROUGE-SU4 shows a better correlation comparing to GROUGE-SU4.

Comparing the best-performing AESOP metrics in Figure 4.2, GROUGE-SU4 achieves the best correlation with Responsiveness when measuring with the Pearson correlation. We also observe that GROUGE-2 obtains the best correlation with Responsiveness while measuring with the Spearman and Kendall rank correlations. The reason is that semantic interpretation of bigrams is easier, and that of contiguous bigrams is

Metric	Pearson	Spearman	Kendall
ROUGE-1	0.966	0.909	0.747
ROUGE-2	0.961	0.894	0.745
ROUGE-SU4	0.981	0.894	0.737
ROUGE-WE-1	0.949	0.914	0.753
ROUGE-WE-2	0.977	0.898	0.744
ROUGE-WE-SU4	0.978	0.881	0.720
GROUGE-1	0.968	0.916	0.758
GROUGE-2	0.979	0.921	0.768
GROUGE-SU4	0.980	0.901	0.747

Table 4.1: Correlation of the **ROUGE-based metrics** with the manual metric of Pyramid using the correlation metrics of Pearson r , Spearman ρ , and Kendall τ on the TAC 2011 AESOP dataset

much more precise. Regarding Table 4.2, every variant of GROUGE outperforms its corresponding variant in the framework of ROUGE.

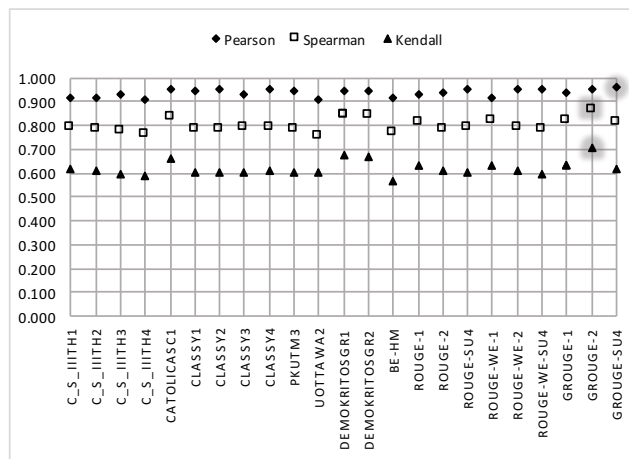


Figure 4.2: Correlation of the **best-performing AESOP metrics** with the manual metric of Responsiveness using the correlation metrics of Pearson r , Spearman ρ , and Kendall τ on the TAC 2011 AESOP dataset

The readability score reflects the fluency and structure of the summary, independently of content; and is based on grammaticality, structure, focus, coherence and etc.. According to Figure 4.3, GROUGE-SU4 and GROUGE-2 are superior to the best-performing AESOP metrics, regarding Pearson and Spearman/Kendall rank correlations, respectively. Although our main goal is not to improve the readability, GROUGE achieves the best correlations with this metric. This is likely due to considering word types and part-of-speech tagging while aligning and disambiguating n-grams. Part-of-speech features are shown by [4] to be helpful in the prediction of the linguistic quality. Measuring Readability in the ROUGE framework, every variant of GROUGE represents the best correlation results comparing to its corresponding variant of ROUGE and ROUGE-WE for all correlation metrics (Table 4.3).

Metric	Pearson	Spearman	Kendall
ROUGE-1	0.935	0.818	0.633
ROUGE-2	0.942	0.790	0.610
ROUGE-SU4	0.955	0.790	0.602
ROUGE-WE-1	0.916	0.819	0.631
ROUGE-WE-2	0.953	0.797	0.615
ROUGE-WE-SU4	0.954	0.787	0.597
GROUGE-1	0.940	0.822	0.635
GROUGE-2	0.954	0.863	0.705
GROUGE-SU4	0.958	0.812	0.617

Table 4.2: Correlation of the **ROUGE-based metrics** with the manual metric of Responsiveness using the correlation metrics of Pearson r , Spearman ρ , and Kendall τ on the TAC 2011 AESOP dataset

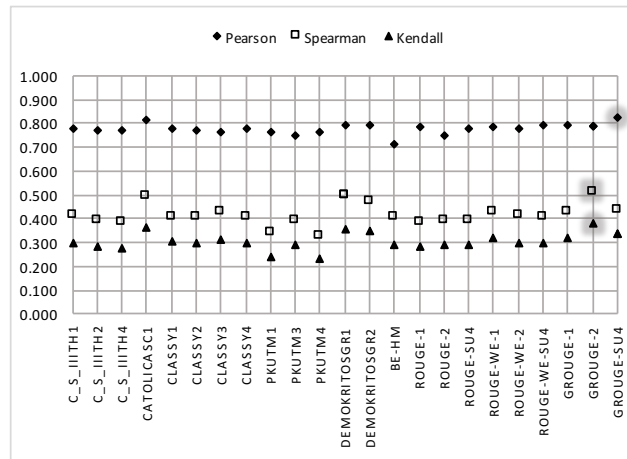


Figure 4.3: Correlation of the **best-performing AESOP metrics** with the manual metric of Readability using the correlation metrics of Pearson r , Spearman ρ , and Kendall τ on the TAC 2011 AESOP dataset

Overall, considering Pyramid, Responsiveness, and Readability, and regardless of the correlation metric used, every GROUGE variant outperforms its corresponding ROUGE variant, with only one exception: ROUGE-SU4 correlates slightly better with Pyramid when measuring with Pearson correlation, to which possible reasons are discussed earlier. Looking at GROUGE-2 that is far more superior than its corresponding variants while measuring with Spearman and Kendall rank correlations, supports our proposal to consider semantics besides surface with ROUGE. However, the large/small differences in competing correlations with human assessment are not an acceptable proof of superiority/inferiority in performance of one metric over another. Hence, prior to any conclusion in this regard, significance tests should be applied.

Metric	Pearson	Spearman	Kendall
ROUGE-1	0.790	0.391	0.285
ROUGE-2	0.752	0.398	0.293
ROUGE-SU4	0.784	0.395	0.293
ROUGE-WE-1	0.785	0.431	0.322
ROUGE-WE-2	0.782	0.414	0.304
ROUGE-WE-SU4	0.793	0.407	0.302
GROUGE-1	0.793	0.433	0.326
GROUGE-2	0.787	0.513	0.384
GROUGE-SU4	0.824	0.440	0.334

Table 4.3: Correlation of the **ROUGE-based metrics** with the manual metric of Readability using the correlation metrics of Pearson r , Spearman ρ , and Kendall τ on the TAC 2011 AESOP dataset

4.3 Significance Test

Evaluation of summarization metrics depart from correlation with human judgment has included the ability of a metric/significance test combination to identify a significant difference between the quality of human and system-generated summaries [21]. To better clarify the effectiveness of GROUGE, we use pairwise Williams significance test² recommended by [6] for summarization evaluation. Accordingly, evaluation of a given summarization metric, M_{new} , takes the form of quantifying three correlations: $r(M_{new}, H)$, that exists between the evaluation metric scores for summarization systems and corresponding human assessment scores; $r(M_{base}, H)$, that stands for the correlation of baseline metrics with human judges; and the third correlation, between evaluation metric scores themselves, $r(M_{base}, M_{new})$. It can happen for a pair of competing metrics for which the correlation between metric scores is strong, that a small difference in competing correlations with human assessment is significant, while, for a different pair of metrics with a larger difference in correlation, the difference is not significant [6]. Utilizing this significance test, the results show that all increases in correlations of GROUGE compared to ROUGE and ROUGE-WE variants in Tables 4.1, 4.2 and 4.3 are statistically significant ($p < 0.05$).

4.4 Exploring Scaling Factor

In this section, we optimize scaling factor β in Equation 3.4, and obtain a balance between contributions of lexical and semantic similarity scores to calculate the lexico-semantic similarity. To this end, we make use of the TAC 2010 AESOP dataset. Figure 4.4 shows the correlation results obtained by the variants of GROUGE with Pyramid (Pyr) and Responsiveness (Rsp) metrics measured by Pearson. The best results are observed when using $\beta = 0.5$. Performance deteriorates when the value of β approaches 1.0 which indicates the ROUGE scores without any touch of semantic similarity. Decreasing the weight of β to zero causes the exclusion of lexical match counts, and consequently inappropriateness of the outcomes. This demonstrates the importance of using both lexical and semantic similarities to fairly judge the quality of summaries.

²Also known as Hotelling-Williams




Figure 4.4: Exploring scaling factor β on the TAC 2010 AESOP dataset

5 Conclusion

We have proposed an effective approach (namely GROUGE) to overcome the limitation of high lexical dependency in ROUGE. We improve on ROUGE by performing both semantic and lexical analysis of summaries. Evaluation is processed by comparing each model-gram against the corresponding peer summary text. To this end, the PPR algorithm is employed, and all senses have been disambiguated before comparison. Experiments over the TAC AESOP datasets demonstrate that GROUGE achieves higher correlations with manual judgments in comparison with the well-established ROUGE. Since this approach goes beyond the lexical surface and exploits the underlying semantics, we believe that it would work even better on more comprehensive texts such as a dataset provided for the evaluation of abstractive summaries. Therefore, our ongoing work includes constructing a standard dataset for assessing the automatic metrics specified to evaluate abstractive summaries. We also believe that this approach can open a door to the evaluation of automatic text simplification. The reason is that text simplification indicates the process of simplifying a text without losing its meaning, and this approach can capture the underlying meaning in a text, regardless of its surface. Hence, in future, we intend to adopt this approach with the aim of helping ROUGE to gain qualitative insights into the nature of text simplification.

Bibliography

- [1] Marco Baroni, Georgiana Dinu, and Germán Kruszewski. Don't count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *ACL (1)*, pages 238–247, 2014.
- [2] Samira Ellouze, Maher Jaoua, and Lamia Hadrich Belguith. An evaluation summary method based on a combination of content and linguistic metrics. In *Recent Advances in Natural Language Processing (RANLP)*, pages 245–251. Citeseer, 2013.
- [3] Christiane Fellbaum. *WordNet*. Wiley Online Library, 1998.
- [4] Lijun Feng, Martin Jansche, Matt Huenerfauth, and Noémie Elhadad. A comparison of features for automatic readability assessment. In *Proceedings of the 23rd*

- International Conference on Computational Linguistics: Posters*, pages 276–284. Association for Computational Linguistics, 2010.
- [5] George Giannakopoulos, Vangelis Karkaletsis, George Vouros, and Panagiotis Stamatopoulos. Summarization system evaluation revisited: N-gram graphs. *ACM Transactions on Speech and Language Processing (TSLP)*, 5(3):5, 2008.
 - [6] Yvette Graham et al. Re-evaluating automatic summarization with bleu and 192 shades of rouge. In *Proceedings of International Conference on Empirical Methods on Natural Language Processing (EMNLP)*, pages 128–137, 2015.
 - [7] Taher H Haveliwala. Topic-sensitive pagerank. In *Proceedings of the 11th international conference on World Wide Web (WWW)*, pages 517–526. ACM, 2002.
 - [8] Eduard Hovy, Chin-Yew Lin, Liang Zhou, and Junichi Fukumoto. Automated summarization evaluation with basic elements. In *Proceedings of the Fifth Conference on Language Resources and Evaluation (LREC 2006)*, pages 604–611. Citeseer, 2006.
 - [9] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the Association for Computational Linguistics (ACL) workshop*, volume 8. Barcelona, Spain, 2004.
 - [10] Chin-Yew Lin and FJ Och. Looking for a few good metrics: Rouge and its evaluation. In *NII Testbeds and Community for Information access Research (NTCIR) Workshop*, 2004.
 - [11] Inderjeet Mani. *Automatic summarization*, volume 3. John Benjamins Publishing, 2001.
 - [12] Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. Linguistic regularities in continuous space word representations. In *Hlt-naacl*, volume 13, pages 746–751, 2013.
 - [13] Jun-Ping Ng and Viktoria Abrecht. Better summarization evaluation with word embeddings for rouge. *arXiv preprint arXiv:1508.06034*, 2015.
 - [14] Paul Over and J Yen. An introduction to duc 2004 intrinsic evaluation of generic new text summarization systems, 2004. *National Institute of Standards and Technology (NIST)*, 2004.
 - [15] Karolina Owczarzak and Hoa Trang Dang. Overview of the tac 2010 summarization track. In *Proceedings of the Third Text Analysis Conference (TAC), Gaithersburg, Maryland, USA. National Institute of Standards and Technology. <http://www.nist.gov/tac/publications>*, 2010.
 - [16] Karolina Owczarzak and Hoa Trang Dang. Overview of the tac 2011 summarization track: Guided task and aesop task. In *Proceedings of the Text Analysis Conference (TAC 2011), Gaithersburg, Maryland, USA, November, 2011*.
 - [17] Rebecca J Passonneau, Emily Chen, Weiwei Guo, and Dolores Perin. Automated pyramid scoring of summaries using distributional semantics. 2013.

- [18] Rebecca J Passonneau, Ani Nenkova, Kathleen McKeown, and Sergey Sigelman. Applying the pyramid method in duc 2005. In *Proceedings of the Document Understanding Conference (DUC 05), Vancouver, BC, Canada, 2005*.
- [19] Mohammad Taher Pilehvar, David Jurgens, and Roberto Navigli. Align, disambiguate and walk: A unified approach for measuring semantic similarity. In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 1341–1351, 2013.
- [20] Mohammad Taher Pilehvar and Roberto Navigli. From senses to texts: An all-in-one graph-based approach for measuring semantic similarity. *Artificial Intelligence*, 228:95–128, 2015.
- [21] Peter Rankel, John M Conroy, Eric V Slud, and Dianne P O’Leary. Ranking human and machine summarization systems. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 467–473. Association for Computational Linguistics, 2011.
- [22] Josef Steinberger and Karel Ježek. Evaluation measures for text summarization. *COMPUTING AND INFORMATICS*, 28(2):251–275, 2012.
- [23] Stephen Tratz and Eduard Hovy. Bewte: basic elements with transformations for evaluation. In *Proceedings of the Text Analysis Conference (TAC) 2008 Workshop*, 2008.
- [24] Joseph Turian, Lev Ratinov, and Yoshua Bengio. Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pages 384–394. Association for Computational Linguistics, 2010.