

Using Statistical and Semantic Models for Multi-Document Summarization

Divyanshu Daiya

Department of Computer Science Engineering

LNM Institute of Information Technology

Jaipur, INDIA

daiyadivyanshu@gmail.com

Anukarsh Singh

Department of Computer Science Engineering

LNM Institute of Information Technology

Jaipur, INDIA

anukarshsingh1@gmail.com

Abstract. We report a series of experiments with different semantic models on top of various statistical models for extractive text summarization. Though statistical models may better capture word co-occurrences and distribution around the text, they fail to detect the context and the sense of sentences /words as a whole. Semantic models help us gain better insight into the context of sentences. We show that how tuning weights between different models can help us achieve significant results on various benchmarks. Learning pre-trained vectors used in semantic models further, on given corpus, can give addition spike in performance. Using weighing techniques in between various statistical models too further refines our result. For Statistical models, we have used TF/IDF, TextRANK, Jaccard/Cosine Similarities. For Semantic Models, we have used WordNet-based Model and proposed two models based on Glove Vectors and Facebook's InferSent. We tested our approach on DUC 2004 dataset, generating 100-word summaries. We have discussed the system, algorithms, analysis and also proposed and tested possible improvements. ROUGE scores were used to compare to other summarizers.

Keywords. Extractive Text Summarization, Semantic Summarization Models, Statistical Summarization Models, Multi Document Summarization

1. Introduction

Automatic Text Summarization deals with the task of condensing documents into a summary, whose level is similar to a human-generated summary. It is mostly distributed into two distinct domains, i.e., Abstractive Summarization and Extractive Summarization. Abstractive summarization(Dejong et al. ,1978) involves models to deduce the crux of the document. It then presents a summary consisting of words and phrases that were not there in the actual document, sometimes even paraphrasing [20]. A state of art method proposed by Wenyuan Zeng [25] produces such summaries with length restricted to 75. There have been many recent developments that produce optimal results, but it is still in a developing phase. It highly relies on natural language processing techniques, which is still evolving to match human standards. These shortcomings make abstractive summarization highly domain selective. As a result, their application is skewed to the areas where NLP techniques have been superlative. Extractive Summarization, on the other hand, uses different methods to identify the most informative/dominant sentences through the text, and then present the results, ranking them accordingly. In this paper, we have proposed two novel stand-alone summarization methods.The first method is based on Glove Model [17],and other is based on Facebook’s InferSent [4]. We have also discussed how we can effectively subdue shortcomings of one model by using it in coalition with models which capture the view that other faintly held.

2. Related Work

A vast number of methods have been used for document summarization. Some of the methods include determining the length and positioning of sentences in the text [19], deducing centroid terms to find the importance of text [19] and setting a threshold on average TF-IDF scores. Bag-of-words approach, i.e., making sentence/Word freq matrix, using a signature set of words and assigning them weights to use them as a criterion for importance measure [10] have also been used. Summarization using weights on high-frequency words [14] describes that high-frequency terms can be used to deduce the core of document.

While semantic summarizers like Lexical similarity is based on the assumption that important sentences are identified by strong chains [6, 2, 13]. In other words, it relates sentences that employ words with the same meaning (synonyms) or other semantic relation. It uses WordNet [12] to find similarity among words that apply to Word Frequency algorithm.POS(Part of Speech) Tagging and WSD(Word Sense Disambiguation) are common among semantic summarizers. Graphical summarizers like TextRank have also provided great benchmark results.TextRank assigns weights to important keywords from the document using graph-based model and sentences which capture most of those concepts/keywords are ranked higher) [2, 11] TextRank

uses Google's PageRank (Brin and Page, 1998) for graphical modeling. Though semantic and graphical models may better capture the sense of document but miss out on statistical view. There is a void of hybrid summarizers; there haven't been many studies made in the area. Wong [23] conducted some preliminary research but there isn't much there on benchmark tests to our knowledge. We use a mixture of statistical and semantic models, assign weights among them by training on field-specific corpora. As there is a significant variation in choices among different fields. We support our proposal with expectations that shortcomings posed by one model can be filled with positives from others. We deploy experimental analysis to test our proposition.

3. Proposed Approach

For Statistical analysis we use Similarity matrices, word co-occurrence/ n-gram model, and TF/IDF matrix. For semantic analysis we use custom Glove based model, WordNet based Model and Facebook InferSent [4] based Model. For Multi-Document Summarization, after training on corpus, we assign weights among the different techniques. We store the sense vector for documents, along with weights, for future reference. For Single document summarization, firstly we calculate the sense vector for that document and calculate the nearest vector from the stored Vectors, we use the weights of the nearest vector. We will describe the flow for semantic and statistical models separately.

3.1. Prepossessing

We discuss, in detail, the steps that are common for both statistical and semantic models.

3.1.1. Sentence Tokenizer

We use NLTK sentence tokenizer `sent_tokenize()`, based on PUNKT tokenizer, pre-trained on a corpus. It can differentiate between Mr. , Mrs. and other abbreviations etc. and the normal sentence boundaries. [7]

Given a document D we tokenize it into sentences as $\langle s_1, s_2, s_3, s_4 \dots s_n \rangle$.

3.1.2. Word Tokenizer

Replacing all the special characters with spaces for easier word-tagging and Tokenizing. We use NLTK word tokenizer, which is a Penn Treebank-style tokenizer, to tokenize words. We calculate the total unique words in the Document. If we can write any sentence as:-

$$s_i \rightarrow \langle \mathbf{w}_I, \mathbf{w}_J, \mathbf{w}_K, \mathbf{w}_L, \dots \rangle, i \in (1, n)$$

Then the number of unique words can be represented as:-

$$(I, J, K, L, \dots) \subset (1, \dots, M), n \rightarrow \text{Totalsentences}, M \rightarrow \text{Totaluniquewords}$$

3.2. Using Stastical Models

3.2.1. Similarity/Correlation Matrices

Frequency Matrix generation: Our tokenized words contain redundancy due to digits and transitional words such as “and”, “but” etc., which carry little information. Such words are termed stop words. [22] We removed stop words and words occurring in $<0.2\%$ and $>15\%$ of the documents (considering the word frequency over all documents). After the removal, the no. of unique words left in the particular document be p where $p < m$ (where m is the total no. of unique words in our tokenized list originally). We now formulate a matrix $F_{n \times p}$ where n is the total number of sentences and p is the total number of unique words left in the document. Element e_{ij} in the matrix $F_{n \times p}$ denotes frequency of j^{th} unique word in the i^{th} sentence.

Similarity/Correlation Matrix generation: We now have have sentence word frequency vector \mathbf{Sf}_i as $\langle \mathbf{f}_{i1}, \mathbf{f}_{i2}, \mathbf{f}_{i3}, \dots, \mathbf{f}_{i4} \rangle$ where f_{ia} denotes frequency of a^{th} unique word in the i^{th} sentence. We now compute,

$$\text{Sentence_similarity}(\mathbf{Sf}_i, \mathbf{Sf}_j)$$

We use two similarity measures : We generate the similarity matrix $Sim_{n \times n}^j$ for each of the similarity Measure, where j indexes the similarity Measure. Element E_{ij} of $Sim_{n \times n}^j$ denotes similarity between i^{th} and j^{th} sentence. Consequentially, we will end up with $Sim_{n \times n}^1$ and $Sim_{n \times n}^2$, corresponding to each similarity measure.

1. Jaccard Similarity: For some sets A and B , $\langle a, b, c, \dots \rangle$ and $\langle x, y, z, \dots \rangle$ respectively, the Jaccard Similarity is defined as:-

$$\text{Jaccard_similarity}(A, B) \leftarrow \frac{n(A \cap B)}{n(A \cup B)}$$

2. Cosine Similarity: The Cosine distance between ‘ u ’ and ‘ v ’, is defined as:-

$$\text{Cosine_similarity}(A, B) \leftarrow 1 - \frac{u \cdot v}{\|u\| \|v\|}$$

where ‘ $u \cdot v$ ’ is the dot product of ‘ u ’ and ‘ v ’.

3.2.2. PageRank

PageRank algorithm [15], devised to rank web pages, forms the core of Google Search. It roughly works by ranking pages according to the number and quality of outsourcing links from the page. For NLP, a PageRank based technique, TextRank has been a major breakthrough in the field. TextRank based summarization has seeded exemplary results on benchmarks. We use a naive TextRank analogous for our task.

Given n sentences $\langle s_1, s_2, s_3, \dots, s_n \rangle$, we intend to generate PageRank or probability distribution matrix $\mathbf{R}_{n \times 1}$,

$$\begin{bmatrix} Pr(s_1) \\ Pr(s_2) \\ \vdots \\ Pr(s_n) \end{bmatrix}$$

, where $Pr(s_k)$ in original paper denoted probability with which a randomly browsing user lands on a particular page. For the summarization task, they denote how strongly a sentence is connected with rest of document, or how well sentence captures multiple views/concepts. The steps are as:

1. Initialize \mathbf{R} as,

$$\begin{bmatrix} Pr(s_1) \\ Pr(s_2) \\ \vdots \\ Pr(s_n) \end{bmatrix} = \begin{bmatrix} \frac{1}{n} \\ \frac{1}{n} \\ \vdots \\ \frac{1}{n} \end{bmatrix}$$

2. Define \mathbf{d} , probability that randomly chosen sentence is in summary and ϵ as measure of change i.e. to stop computation when difference between successive \mathbf{R} computations recedes below ϵ .
3. Using cosine-similarity matrix $Sim_{n \times n}^2$, we generate the following equation as a measure for relation between sentences:-

$$\mathbf{R} = \begin{bmatrix} (1-d)/n \\ (1-d)/n \\ \vdots \\ (1-d)/n \end{bmatrix} + \mathbf{d} \times \mathbf{Sim}_{n \times n}^2 \times \mathbf{R}$$

4. Repeat last step until $|R(t+1) - R(t)| > \epsilon$.
5. Take top ranking sentences in \mathbf{R} for summary.

3.2.3. TF/IDF

Term Frequency(TF)/Bag of words is the count of how many times a word occurs in the given document. Inverse Document Frequency(IDF) is the number of times word occurs in complete corpus. Infrequent words through corpus will have higher weights, while weights for more frequent words will be deprecated.

Underlying steps for TF/IDF summarization are:

1. Create a count vector

$$Doc_1 \leftarrow \langle fr_{Word_1}, fr_{Word_2}, fr_{Word_3}, \dots \rangle$$

2. Build a tf-idf matrix $W_{M \times N}$ with element $w_{i,j}$ as,

$$w_{i,j} = tf_{i,j} \times \log\left(\frac{N}{df_i}\right)$$

Here, $tf_{i,j}$ denotes term frequency of i th word in j th sentence, and $\log\left(\frac{N}{df_i}\right)$ represents the IDF frequency.

3. Score each sentence, taking into consideration only nouns, we use NLTK POS-tagger for identifying nouns.

$$Score(S_{o,j}) \leftarrow \frac{\sum N_{o,j}}{\sum_{p=1}^N N_{p,j}}$$

4. Applying positional weighing .

$$Scores[S_{o,j}] = Score(S_{o,j}) \times \left(\frac{o}{T}\right)$$

$o \rightarrow$ Sentence index $T \rightarrow$ Total sentences in document j

5. Summarize using top ranking sentences.

3.3. Using Semantic Models

We proceed in the same way as we did for statistical models. All the pre-processing steps remain nearly same. We can make a little change by using lemmatizer instead of stemmer. Stemming involves removing the derivational affixes/end of words by heuristic analysis in hope to achieve base form. Lemmatization, on the other hand, involves firstly POS tagging [21], and after morphological and vocabulary analysis, reducing the word to its base form. Stemmer output for 'goes' is 'goe', while lemmatized output with the verb passed as POS tag is 'go'. Though lemmatization may have little more time overhead as compared to stemming, it necessarily provides better base word reductions. Since WordNet [16] and Glove both require dictionary look-ups, in order for them to work well, we need better base word mappings. Hence lemmatization is preferred.

3.3.1. Additional Pre-processing

1. **Part of Speech(POS) Tagging:** We tag the words using NLTK POS-Tagger.
2. **Lemmatization:** We use NTLK lemmatizer with POS tags passed as contexts.

3.3.2. Using WordNet

We generated Similarity matrices in the case of Statistical Models. We will do the same here, but for sentence similarity measure we use the method devised by Dao. [5] The method is defined as:

1. **Word Sense Disambiguation(WSD):** We use the adapted version of Lesk algorithm [9], as devised by Dao, to derive the sense for each word.
2. **Sentence pair Similarity:** For each pair of sentences, we create semantic similarity matrix S . Let A and B be two sentences of lengths m and n respectively. Then the resultant matrix S will be of size $m \times n$, with element $s_{i,j}$ denoting semantic similarity between sense/synset of word at position i in sentence A and sense/synset of word at position j in sentence B , which is calculated by path length similarity using `is-a` (hypernym/hyponym) hierarchies. It uses the idea that shorter the path length, higher the similarity. To calculate the path length, we proceed in following manner:-

For two words W_1 and W_2 , with synsets s_1 and s_2 respectively,

$$sd(s_1, s_2) = 1 / distance(s_1, s_2)$$

$$S_{m \times n} = \begin{bmatrix} sd(s_1, s_1) & \dots & sd(s_1, s_n) \\ sd(s_2, s_1) & \ddots & \vdots \\ \vdots & sd(s_i, s_j) & \\ sd(s_m, s_1) & \dots & sd(s_m, s_n) \end{bmatrix}$$

We formulate the problem of capturing semantic similarity between sentences as the problem of computing a maximum total matching weight of a bipartite graph, where X and Y are two sets of disjoint nodes. We use the Hungarian method [8] to solve this problem. Finally we get bipartite matching matrix B with entry $b_{i,j}$ denoting matching between $A[i]$ and $B[j]$. To obtain the overall similarity, we use Dice coefficient,

$$Sim(A, B) = \frac{|A \cap B|}{|A| + |B|}$$

with threshold set to 0.5, and $|A|, |B|$ denoting lengths of sentence A and B respectively.

3. We perform the previous step over all pairs to generate the similarity matrix $Sim_{N \times N}^3$.

3.3.3. Using Glove Model

Glove Model provides us with a convenient method to represent words as vectors, using vectors representation for words, we generate vector representation for sentences. We work in the following order,

1. Represent each tokenized word w_i in its vector form $\langle a_i^1, a_i^2, a_i^3, \dots, a_i^{300} \rangle$.
2. Represent each sentence into vector using following equation,

$$SVec(s_j) = \frac{1}{|s_j|} \sum_{w_i \in s_j} f_{i,j}(a_i^1, a_i^2, \dots, a_i^{300})$$

where $f_{i,j}$ being frequency of w_i in s_j .

3. Calculate similarity between sentences using cosine distance between two sentence vectors.
4. Populate similarity matrix $Sim_{N \times N}^4$ using previous step.

3.3.4. Using Facebook's InferSent

InferSent is a state of the art supervised sentence encoding technique [4]. It outperformed another state-of-the-art sentence encoder SkipThought on several benchmarks, like the STS benchmark (<http://ixa2.si.ehu.es/stswiki/index.php/STSbenchmark>). The model is trained on Stanford Natural Language Inference (SNLI) dataset [3] using seven architectures Long Short-Term Memory (LSTM), Gated Recurrent Units (GRU), forward and backward GRU with hidden states concatenated, Bi-directional LSTMs (BiLSTM) with min/max pooling, self-attentive network and (HCN's) Hierarchical convolutional networks. The network performances are task/corpus specific.

Steps to generate similarity matrix $Sim_{N \times N}^5$ are:

1. Encode each sentence to generate its vector representation $\langle l_i^1, l_i^2, l_i^3, \dots, l_i^{4096} \rangle$.
2. Calculate similarity between sentence pair using cosine distance.
3. Populate similarity matrix $Sim_{N \times N}^5$ using previous step.

3.4. Generating Summaries

TF-IDF scores and TextRank allows us to directly rank sentences and choose k top sentences, where k is how many sentences user want in the summary. On the other hand, the similarity matrix based approach is used in case of all Semantic Models, and Similarity/correlation based Statistical models. To rank sentences from Similarity matrix, we can use following approaches:-

1. Ranking through Relevance score

For each sentence s_i in similarity matrix the Relevance Score is as:-

$$RScore(s_i) = \sum_{j=1}^N Sim[i, j]$$

We can now choose k top ranking sentences by RScores. Higher the RScore, higher the rank of sentence.

2. Hierarchical Clustering

Given a similarity matrix $Sim_{N \times N}$, let $s_{a,b}$ denote an individual element, then Hierarchical clustering is performed as follows:-

- (a) Initialize a empty list R .
- (b) Choose element with highest similarity value let it be $s_{i,j}$ where, $i \neq j, s_{i,j} \neq 0$
- (c) Replace values in column and row i in following manner:-

$$s_{d,i} = \frac{s_{d,i} + s_{d,j}}{2}, d \in (1, N), \quad s_{i,d} = \frac{s_{i,d} + s_{j,d}}{2}, d \in (1, N)$$

- (d) Replace entries corresponding to column and row i by zeros.
- (e) Add i and j to R , if they are not already there.
- (f) Repeat steps 2-5 until single single non-zero element remains, for remaining non-zero element apply Step 5 and terminate.
- (g) We will have rank list R in the end.

We can now choose k top ranking sentences from R .

3.5. Single Document Summarization

After generating summary from a particular model, our aim is to compute summaries through overlap of different models. Let us have g summaries from g different models. For p_{th} summarization model, let the k sentences contained be:-

$$Sum_p \leftarrow (s_{(1,p)}, s_{(2,p)} \cdots, s_{(k,p)})$$

Now for our list of sentences $\langle s_1, s_2, s_3, \dots, s_n \rangle$ we define **cWeight** as weight obtained for each sentence using g models.

$$cWeight(s_i) = \sum_{j=1}^g W_j B(j, s_i)$$

Here, $B(j, s_i)$ is a function which returns 1 if sentence is in summary of j_{th} model, otherwise zero. W_i is weight assigned to each model without training, $W_i = \frac{1}{g}, i \in (1, g)$

3.6. Multi-Document/Domain-Specific Summarization

We here use machine learning based approach to further increase the quality of our summarization technique. The elemental concept is that we use training set of u domain specific documents, with gold standard/human-composed summaries, provided we fine tune our weights $W_i \forall i \in (1, g)$ for different models taking F1-score/F-measure. [18] as factor.

$$F1Score = \frac{2 \cdot precision \cdot recall}{precision + recall}$$

We proceed in the following manner:-

1. For each document in training set generate summary using each model independently, compute the $F1Score$ w.r.t. gold summary.
2. For each model, assign the weights using

$$W_i = \frac{\sum_{j=1}^v f_1^{(j,i)}}{u}, i \in (1, g)$$

Here, $f_1^{(j,i)}$ denotes $F1Score$ for j_{th} model in i_{th} document.

We now obtain **cWeight** as we did previously, and formulate cumulative summary, capturing the consensus of different models. We hence used a supervised learning algorithm to capture the mean performances of different models over the training data to fine-tune our summary.

3.7. Domain-Specific Single Document Summarization

As we discussed earlier, summarization models are field selective. Some models tend to perform remarkably better than others in certain fields. So, instead of assigning uniform weights to all models we can go by the following approach.

1. For each set of documents we train on, we generate document vector using bidirectional GRU ([1] as described by Zichao Yang [24]) for each document. We then generate complete corpus vector as follows:-

$$cDocs = \sum_{i=1}^v (a_i^1, a_i^1, a_i^1, \dots, a_i^p,)$$

where, v is total training set size, p is number of features in document vector.

2. We save $cDocs$ and $weights$ corresponding to each corpus.
3. For each single document summarization task, we generate given texts document vector, perform nearest vector search over all stored $cDocs$, apply weights corresponding to that corpus.

3.8. Experiments

Table 1: Average ROUGE-2 Scores for Different Combination of Models.

Models						Score
A	B	C	D	E	F	<i>ROUGE2(95%)</i>
		•		•		0.03172
•				•		0.03357
•		•		•		0.03384
	•	•		•		0.03479
•	•			•		0.03572
•	•	•		•		0.03519
		•			•	0.03821
•					•	0.03912
•		•			•	0.03822
	•	•			•	0.03986
•	•				•	0.04003
•	•	•			•	0.03846
		•	•			0.03312
•			•			0.03339
•		•	•			0.03332
	•	•	•			0.03532
•	•		•			0.03525
•	•	•	•			0.03519
		•	•		•	0.03721
•			•		•	0.03689
•		•	•		•	0.03771
	•	•	•		•	0.03812
•	•		•		•	0.03839
•	•	•	•		•	0.03782
		•		•	•	0.03615
•				•	•	0.03598
•		•		•	•	0.03621
	•	•		•	•	0.03803

Continued on next page

Table 1 – Continued from previous page

Models						Score
A	B	C	D	E	F	<i>ROUGE2(95%)</i>
•	•			•	•	0.03819
•	•	•		•	•	0.03784
		•	•	•		0.03314
•			•	•		0.03212
•		•	•	•		0.03426
	•	•	•	•		0.03531
•	•		•	•		0.03544
•	•	•	•	•		0.03529
		•	•	•	•	0.03712
•			•	•	•	0.03713
•		•	•	•	•	0.03705
	•	•	•	•	•	0.03821
•	•		•	•	•	0.03829
•	•	•	•	•	•	0.03772

We evaluate our approaches on 2004 DUC(Document Understanding Conferences) dataset(<https://duc.nist.gov/>). The Dataset has 5 Tasks in total. We work on Task 2. It (Task 2) contains 50 news documents cluster for multi-document summarization. Only 665-character summaries are provided for each cluster. For evaluation, we use ROUGE, an automatic summary evaluation metric. It was firstly used for DUC 2004 data-set. Now, it has become a benchmark for evaluation of automated summaries. ROUGE is a correlation metric for fixed-length summaries populated using n-gram co-occurrence. For comparison between model summary and to-be evaluated summary, separate scores for 1, 2, 3, and 4-gram matching are kept. We use ROUGE-2, a bi-gram based matching technique for our task.

In the *Table 1*, we try different model pairs with weights trained on corpus for Task 2. We have displayed mean ROUGE-2 scores for base Models. We have calculated final scores taking into consideration all normalizations, stemming, lemmatizing and clustering techniques, and the ones providing best results were used. We generally expected WordNet, Glove based semantic models to perform better given they better capture crux of the sentence and compute

¹**A**→ Jaccard/Cosine Similarity Matrix **B**→ TextRank **C**→ TFIDF **D**→ WordNet Based Model **E**→ Glove-vec Based Model **F**→ InferSent Based Model

similarity using the same, but instead, they performed average. This is attributed to the fact they assigned high similarity scores to not so semantically related sentences. We also observe that combinations with TF/IDF and Similarity Matrices(Jaccard/Cosine) offer nearly same results. The InferSent based Summarizer performed exceptionally well. We initially used pre-trained features to generate sentence vectors through InferSent.

Table 2: Average ROUGE-2 scores for base methods.

Model	<i>ROUGE – 2</i>
Jaccard	0.03468
Cosine	0.02918
TextRank	0.03629
TFIDF	0.03371
WordNet Based Model	0.03354
Glove-vec Based Model	0.03054
InferSent Based Model	0.03812

3.9. Conclusion/Future Work

We can see that using a mixture of Semantic and Statistical models offers an improvement over stand-alone models. Given better training data, results can be further improved. Using domain-specific labeled data can provide a further increase in performances of Glove and WordNet Models.

Some easy additions that can be worked on are:

1. Unnecessary parts of the sentence can be trimmed to improve summary further.
2. Using better algorithm to capture sentence vector through Glove Model can improve results.
3. Query specific summarizer can be implemented with little additions.
4. For generating summary through model overlaps, we can also try Graph-based methods or different Clustering techniques.

References

- [1] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [2] A. Barrera and R. Verma, “Combining syntax and semantics for automatic extractive single-document summarization,” in *International Conference on Intelligent Text Processing and Computational Linguistics*. Springer, 2012, pp. 366–377.

- [3] S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning, “A large annotated corpus for learning natural language inference,” *arXiv preprint arXiv:1508.05326*, 2015.
- [4] A. Conneau, D. Kiela, H. Schwenk, L. Barrault, and A. Bordes, “Supervised learning of universal sentence representations from natural language inference data,” *arXiv preprint arXiv:1705.02364*, 2017.
- [5] T. N. Dao and T. Simpson, “Measuring similarity between sentences,” *The Code Project*, 2005.
- [6] P. Gupta, V. S. Pendluri, and I. Vats, “Summarizing text by ranking text units according to shallow linguistic features,” in *Advanced communication technology (ICACT), 2011 13th international conference on*. IEEE, 2011, pp. 1620–1625.
- [7] T. Kiss and J. Strunk, “Unsupervised multilingual sentence boundary detection,” *Computational Linguistics*, vol. 32, no. 4, pp. 485–525, 2006.
- [8] H. W. Kuhn, “The hungarian method for the assignment problem,” *Naval Research Logistics (NRL)*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [9] M. Lesk, “Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone,” in *Proceedings of the 5th annual international conference on Systems documentation*. ACM, 1986, pp. 24–26.
- [10] C.-Y. Lin and E. Hovy, “The automated acquisition of topic signatures for text summarization,” in *Proceedings of the 18th conference on Computational linguistics-Volume 1*. Association for Computational Linguistics, 2000, pp. 495–501.
- [11] R. Mihalcea and P. Tarau, “Textrank: Bringing order into text,” in *Proceedings of the 2004 conference on empirical methods in natural language processing*, 2004.
- [12] G. A. Miller, R. Beckwith, C. Fellbaum, D. Gross, and K. J. Miller, “Introduction to wordnet: An on-line lexical database,” *International journal of lexicography*, vol. 3, no. 4, pp. 235–244, 1990.
- [13] V. G. Murdock, “Aspects of sentence retrieval,” MASSACHUSETTS UNIV AMHERST DEPT OF COMPUTER SCIENCE, Tech. Rep., 2006.
- [14] A. Nenkova, L. Vanderwende, and K. McKeown, “A compositional context sensitive multi-document summarizer: exploring the factors that influence summarization,” in *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2006, pp. 573–580.
- [15] L. Page, S. Brin, R. Motwani, and T. Winograd, “The pagerank citation ranking: Bringing order to the web.” Stanford InfoLab, Tech. Rep., 1999.

- [16] T. Pedersen, S. Patwardhan, and J. Michelizzi, “Wordnet:: Similarity: measuring the relatedness of concepts,” in *Demonstration papers at HLT-NAACL 2004*. Association for Computational Linguistics, 2004, pp. 38–41.
- [17] J. Pennington, R. Socher, and C. Manning, “Glove: Global vectors for word representation,” in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [18] D. M. Powers, “Evaluation: from precision, recall and f-measure to roc, informedness, markedness and correlation,” 2011.
- [19] D. R. Radev, H. Jing, M. Styś, and D. Tam, “Centroid-based summarization of multiple documents,” *Information Processing & Management*, vol. 40, no. 6, pp. 919–938, 2004.
- [20] T. Rocktäschel, E. Grefenstette, K. M. Hermann, T. Kociský, and P. Blunsom, “Reasoning about entailment with neural attention,” *CoRR*, vol. abs/1509.06664, 2015. [Online]. Available: <http://arxiv.org/abs/1509.06664>
- [21] B. Santorini, “Part-of-speech tagging guidelines for the penn treebank project (3rd revision),” *Technical Reports (CIS)*, p. 570, 1990.
- [22] W. J. Wilbur and K. Sirotkin, “The automatic identification of stop words,” *Journal of information science*, vol. 18, no. 1, pp. 45–55, 1992.
- [23] K.-F. Wong, M. Wu, and W. Li, “Extractive summarization using supervised and semi-supervised learning,” in *Proceedings of the 22nd International Conference on Computational Linguistics-Volume 1*. Association for Computational Linguistics, 2008, pp. 985–992.
- [24] Z. Yang, D. Yang, C. Dyer, X. He, A. Smola, and E. Hovy, “Hierarchical attention networks for document classification,” in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016, pp. 1480–1489.
- [25] W. Zeng, W. Luo, S. Fidler, and R. Urtasun, “Efficient summarization with read-again and copy mechanism,” *CoRR*, vol. abs/1611.03382, 2016. [Online]. Available: <http://arxiv.org/abs/1611.03382>