

Learning to Create Sentence Semantic Relation Graphs for Multi-Document Summarization

Diego Antognini and Boi Faltings

Artificial Intelligence Laboratory

École Polytechnique Fédérale de Lausanne

Lausanne, Switzerland

{diego.antognini, boi.faltings}@epfl.ch

Abstract

Linking facts across documents is a challenging task, as the language used to express the same information in a sentence can vary significantly, which complicates the task of multi-document summarization. Consequently, existing approaches heavily rely on hand-crafted features, which are domain-dependent and hard to craft, or additional annotated data, which is costly to gather. To overcome these limitations, we present a novel method, which makes use of two types of sentence embeddings : universal embeddings, which are trained on a large unrelated corpus, and domain-specific embeddings, which are learned during training. To this end, we develop *SemSentSum*, a fully data-driven model able to leverage both types of sentence embeddings by building a sentence semantic relation graph. *SemSentSum* achieves competitive results on two types of summary, consisting of 665 bytes and 100 words. Unlike other state-of-the-art models, neither hand-crafted features nor additional annotated data are necessary, and the method is easily adaptable for other tasks. To our knowledge, we are the first to use multiple sentence embeddings for the task of multi-document summarization.

1 Introduction

Today’s increasing flood of information on the web creates a need for automated multi-document summarization systems that produce high quality summaries. However, producing summaries in a multi-document setting is difficult, as the language used to display the same information in a sentence can vary significantly, making it difficult for summarization models to capture. Given the complexity of the task and the lack of datasets, most researchers use extractive summarization, where the final summary is composed of existing sentences in the input documents. More specifically,

extractive summarization systems output summaries in two steps : via sentence ranking, where an importance score is assigned to each sentence, and via the subsequent sentence selection, where the most appropriate sentence is chosen, by considering 1) their importance and 2) their frequency among all documents. Due to data sparsity, models heavily rely on well-designed features at the word level (Hong and Nenkova, 2014; Cao et al., 2015; Christensen et al., 2013; Yasunaga et al., 2017) or take advantage of other large, manually annotated datasets and then apply transfer learning (Cao et al., 2017). Additionally, most of the time, all sentences in the same collection of documents are processed independently and therefore, their relationships are lost.

In realistic scenarios, features are hard to craft, gathering additional annotated data is costly, and the large variety in expressing the same fact cannot be handled by the use of word-based features only, as is often the case. In this paper, we address these obstacles by proposing to simultaneously leverage two types of sentence embeddings, namely embeddings pre-trained on a large corpus that capture a variety of meanings and domain-specific embeddings learned during training. The former is typically trained on an unrelated corpus composed of high quality texts, allowing to cover additional contexts for each encountered word and sentence. Hereby, we build on the assumption that sentence embeddings capture both the syntactic and semantic content of sentences. We hypothesize that using two types of sentence embeddings, general and domain-specific, is beneficial for the task of multi-document summarization, as the former captures the most common semantic structures from a large, general corpus, while the latter captures the aspects related to the domain.

We present *SemSentSum* (Figure 1), a fully data-driven summarization system, which does not de-

pend on hand-crafted features, nor additional data, and is thus domain-independent. It first makes use of general sentence embedding knowledge to build a sentence semantic relation graph that captures sentence similarities (Section 2.1). In a second step, it trains genre-specific sentence embeddings related to the domains of the collection of documents, by utilizing a sentence encoder (Section 2.2). Both representations are afterwards merged, by using a graph convolutional network (Kipf and Welling, 2017) (Section 2.3). Then, it employs a linear layer to project high-level hidden features for individual sentences to salience scores (Section 2.4). Finally, it greedily produces relevant and non-redundant summaries by using sentence embeddings to detect similarities between candidate sentences and the current summary (Section 2.6).

The main contributions of this work are as follows :

- We aggregate two types of sentences embeddings using a graph representation. They share different properties and are consequently complementary. The first one is trained on a large unrelated corpus to model general semantics among sentences, whereas the second is domain-specific to the dataset and learned during training. Together, they enable a model to be domain-independent as it can be applied easily on other domains. Moreover, it could be used for other tasks including detecting information cascades, query-focused summarization, keyphrase extraction and information retrieval.
- We devise a competitive multi-document summarization system, which does not need hand-crafted features nor additional annotated data. Moreover, the results are competitive for 665-byte and 100-word summaries. Usually, models are compared in one of the two settings but not both and thus lack comparability.

2 Method

Let C denote a collection of related documents composed of a set of documents $\{D_i | i \in [1, N]\}$ where N is the number of documents. Moreover, each document D_i consists of a set of sentences $\{S_{i,j} | j \in [1, M]\}$, M being the number of sentences in D_i . Given a collection of related documents C , our goal is to produce a summary Sum using a subset of these in the input do-

cuments ordered in some way, such that $Sum = (S_{i_1, j_1}, S_{i_2, j_2}, \dots, S_{i_n, j_n})$.

In this section, we describe how *SemSentSum* estimates the salience score of each sentence and how it selects a subset of these to create the final summary. The architecture of *SemSentSum* is depicted in Figure 1.

In order to perform sentence selection, we first build our sentence semantic relation graph, where each vertex is a sentence and edges capture the semantic similarity among them. At the same time, each sentence is fed into a recurrent neural network, as a sentence encoder, to generate sentence embeddings using the last hidden states. A single-layer graph convolutional neural network is then applied on top, where the sentence semantic relation graph is the adjacency matrix and the sentence embeddings are the node features. Afterward, a linear layer is used to project high-level hidden features for individual sentences to salience scores, representing how salient a sentence is with respect to the final summary. Finally, based on this, we devise an innovative greedy method that leverages sentence embeddings to detect redundant sentences and select sentences until reaching the summary length limit.

2.1 Sentence Semantic Relation Graph

We model the semantic relationship among sentences using a graph representation. In this graph, each vertex is a sentence $S_{i,j}$ (j 'th sentence of document D_i) from the collection documents C and an undirected edge between S_{i_u, j_u} and S_{i_v, j_v} indicates their degree of similarity. In order to compute the semantic similarity, we use the model of Pagliardini et al. (2018) trained on the English Wikipedia corpus. In this manner, we incorporate general knowledge (i.e. not domain-specific) that will complete the specialized sentence embeddings obtained during training (see Section 2.2). We process sentences by their model and compute the cosine similarity between every sentence pair, resulting in a complete graph. However, having a complete graph alone does not allow the model to leverage the semantic structure across sentences significantly, as every sentence pair is connected, and likewise, a sparse graph does not contain enough information to exploit semantic similarities. Furthermore, all edges have a weight above zero, since it is very unlikely that two sentence embeddings are completely orthogonal. To overcome this




Figure 1: Overview of *SemSentSum*. This illustration includes two documents in the collection, where the first one has three sentences and the second two. A sentence semantic relation graph is firstly built and each sentence node is processed by an encoder network at the same time. Thereafter, a single-layer graph convolutional network is applied on top and produces high-level hidden features for individual sentences. Then, saliency scores are estimated using a linear layer and used to produce the final summary.

problem, we introduce an edge-removal-method, where every edge below a certain threshold t_{sim}^g is removed in order to emphasize high sentence similarity. Nonetheless, t_{sim}^g should not be too large, as we otherwise found the model to be prone to overfitting. After removing edges below t_{sim}^g , our sentence semantic relation graph is used as the adjacency matrix A . The impact of t_{sim}^g with different values is shown in Section 3.7.

Based on our aforementioned hypothesis that a combination of general and genre-specific sentence embeddings is beneficial for the task of multi-document summarization, we further incorporate general sentence embeddings, pre-trained on Wikipedia entries, into edges between sentences. Additionally, we compute specialised sentence embeddings, which are related to the domains of the documents (see Section 3.7).

Note that 1) the pre-trained sentence embeddings are only used to compute the weights of the edges and are not used by the summarization model (as others are produced by the sentence encoder) and 2) the edge weights are static and do not change during training.

2.2 Sentence Encoder

Given a list of documents C , we encode each document’s sentence $S_{i,j}$, where each has at most L words ($w_{i,j,1}, w_{i,j,2}, \dots, w_{i,j,L}$). In our experiments, all words are kept and converted into word embeddings, which are then fed to the sentence encoder in order to compute specialised sentence embeddings $S'_{i,j}$. We employ a single-layer

forward recurrent neural network, using Long Short-Term Memory (LSTM) of (Hochreiter and Schmidhuber, 1997) as sentence encoder, where the sentence embeddings are extracted from the last hidden states. We then concatenate all sentence embeddings into a matrix X which constitutes the input node features that will be used by the graph convolutional network.

2.3 Graph Convolutional Network

After having computed all sentence embeddings and the sentence semantic relation graph, we apply a single-layer Graph Convolutional Network (GCN) from Kipf and Welling (2017), in order to capture high-level hidden features for each sentence, encapsulating sentence information as well as the graph structure.

We believe that our sentence semantic relation graph contains information not present in the data (via universal embeddings) and thus, we leverage this information by running a graph convolution on the first order neighborhood.

The GCN model takes as input the node features matrix X and a squared adjacency matrix A . The former contains all sentence embeddings of the collection of documents, while the latter is our underlying sentence semantic relation graph. It outputs hidden representations for each node that encode both local graph structure and nodes’s features. In order to take into account the sentences themselves during the information propagation, we add self-connections (i.e. the identity matrix) to A such that $\tilde{A} = A + I$.

Subsequently, we obtain our sentence hidden features by using Equation 1.

$$S''_{i,j} = \text{ELU}(\tilde{A} \text{ELU}(\tilde{A}XW_0 + b_0)W_1 + b_1) \quad (1)$$

where W_i is the weight matrix of the i 'th graph convolution layer and b_i the bias vector. We choose the Exponential Linear Unit (ELU) activation function from Clevert et al. (2016) due to its ability to handle the vanishing gradient problem, by pushing the mean unit activations close to zero and consequently facilitating the backpropagation. By using only one hidden layer, as we only have one input-to-hidden layer and one hidden-to-output layer, we limit the information propagation to the first order neighborhood.

2.4 Saliency Estimation

We use a simple linear layer to estimate a saliency score for each sentence and then normalize the scores via softmax and obtain our estimated saliency score $S^s_{i,j}$.

2.5 Training

Our model *SemSentSum* is trained in an end-to-end manner and minimizes the cross-entropy loss of Equation 2 between the saliency score prediction and the ROUGE-1 F_1 score for each sentence.

$$\mathcal{L} = - \sum_C \sum_{D \in C} \sum_{S \in D} F_1(S) \log S^s \quad (2)$$

$F_1(S)$ is computed as the ROUGE-1 F_1 score, unlike the common practice in the area of single and multi-document summarization as recall favors longer sentences whereas F_1 prevents this tendency. The scores are normalized via softmax.

2.6 Summary Generation Process

While our model *SemSentSum* provides estimated saliency scores, we use a greedy strategy to construct an informative and non-redundant summary Sum . We first discard sentences having less than 9 words, as in (Erkan and Radev, 2004), and then sort them in descending order of their estimated saliency scores. We iteratively dequeue the sentence having the highest score and append it to the current summary Sum if it is non-redundant with respect to the current content of Sum . We iterate until reaching the summary length limit.

To determine the similarity of a candidate sentence with the current summary, a sentence is considered as dissimilar if and only if the cosine

similarity between its sentence embeddings and the embeddings of the current summary is below a certain threshold t^s_{sim} . We use the pre-trained model of Pagliardini et al. (2018) to compute sentence as well as summary embeddings, similarly to the sentence semantic relation graph construction. Our approach is novel, since it focuses on the semantic sentence structures and captures similarity between sentence meanings, instead of focusing on word similarities only, like previous TF-IDF approaches ((Hong and Nenkova, 2014; Cao et al., 2015; Yasunaga et al., 2017; Cao et al., 2017)).

3 Experiments

3.1 Datasets

We conduct experiments on the most commonly used datasets for multi-document summarization from the Document Understanding Conferences (DUC).¹ We use DUC 2001, 2002, 2003 and 2004 as the tasks of generic multi-document summarization, because they have been carried out during these years. We use DUC 2001, 2002, 2003 and 2004 for generic multi-document summarization, where DUC 2001/2002 are used for training, DUC 2003 for validation and finally, DUC 2004 for testing, following the common practice.

3.2 Evaluation Metric

For the evaluation, we use ROUGE (Lin, 2004) with the official parameters of the DUC tasks and also truncate the summaries to 100 words for DUC 2001/2002/2003 and to 665 bytes for DUC 2004.² Notably, we take ROUGE-1 and ROUGE-2 recall scores as the main metrics for comparison between produced summaries and golden ones as proposed by (Owczarzak et al., 2012). The goal of the ROUGE-N metric is to compute the ratio of the number of N-grams from the generated summary matching these of the human reference summaries.

3.3 Model Settings

To define the edge weights of our sentence semantic relation graph, we employ the 600-dimensional pre-trained unigram model of Pagliardini et al. (2018), using English Wikipedia as source corpus. We keep only edges having a weight larger than $t^g_{sim} = 0.5$ (tuned on the

1. <https://www-nlpir.nist.gov/projects/duc/guidelines.html>

2. ROUGE-1.5.5 with options: -n 2 -m -u -c 95 -x -r 1000 -f A -p 0.5 -t 0 and -l 100 if using DUC 2001/2002/2003 otherwise -b 665.

validation set). For word embeddings, the 300-dimensional pre-trained GloVe embeddings (Pennington et al., 2014) are used and fixed during training. The output dimension of the sentence embeddings produced by the sentence encoder is the same as that of the word embeddings, i.e. 300. For the graph convolutional network, the number of hidden units is 128 and the size of the generated hidden feature vectors is also 300. We use a batch size of 1, a learning rate of 0.0075 using Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$. In order to make *SemSentSum* generalize better, we use dropout (Srivastava et al., 2014) of 0.2, batch normalization (Ioffe and Szegedy, 2015), clip the gradient norm at 1.0 if higher, add L2-norm regularizer with a regularization factor of 10^{-12} and train using early stopping with a patience of 10 iterations. Finally, the similarity threshold t_{sim}^s in the summary generation process is 0.8 (tuned on the validation set).

3.4 Summarization Performance

We train our model *SemSentSum* on DUC 2001/2002, tune it on DUC 2003 and assess the performance on DUC 2004. In order to fairly compare *SemSentSum* with other models available in the literature, experiments are conducted with summaries truncated to 665 bytes (official summary length in the DUC competition), but also with summaries with a length constraint of 100 words. To the best of our knowledge, we are the first to conduct experiments on both summary lengths and compare our model with other systems producing either 100 words or 665 bytes summaries.

3.5 Sentence Semantic Relation Graph Construction

We investigate different methods to build our sentence semantic relation graph and vary the value of t_{sim}^g from 0.0 to 0.75 to study the impact of the threshold cut-off. Among these are :

1. *Cosine* : Using cosine similarity ;
2. *Tf-idf* : Considering a node as the query and another as document. The weight corresponds to the cosine similarity between the query and the document ;
3. *TextRank* (Mihalcea and Tarau, 2004) : A weighted graph is created where nodes are

sentences and edges defined by a similarity measure based on word overlap. Afterward, an algorithm similar to PageRank (Page et al., 1998) is used to compute sentence importance and refined edge weights ;

4. *LexRank* (Erkan and Radev, 2004) : An unsupervised multi-document summarizer based on the concept of eigenvector centrality in a graph of sentences to set up the edge weights ;
5. *Approximate Discourse Graph* (ADG) (Christensen et al., 2013) : Approximation of a discourse graph where nodes are sentences and edges (S_u, S_v) indicates sentence S_v can be placed after S_u in a coherent summary ;
6. *Personalized ADG* (PADG) (Yasunaga et al., 2017) : Normalized version of ADG where sentence nodes are normalized over all edges.

3.6 Ablation Study

In order to quantify the contribution of the different components of *SemSentSum*, we try variations of our model by removing different modules one at a time. Our two main elements are the sentence encoder (*Sent*) and the graph convolutional neural network (*GCN*). When we omit *Sent*, we substitute it with the pre-trained sentence embeddings used to build our sentence semantic relation graph.

3.7 Results and Discussion

Three dimensions are used to evaluate our model *SemSentSum* : 1) the summarization performance, to assess its capability 2) the impact of the sentence semantic relation graph generation using various methods and different thresholds t_{sim}^g 3) an ablation study to analyze the importance of each component of *SemSentSum*.

Summarization Performance We compare the results of *SemSentSum* for both settings : 665 bytes and 100 words summaries. We only include models using the same parameters to compute the ROUGE-1/ROUGE-2 score and recall as metrics.

The results for 665 bytes summaries are reported in Table 1. We compare *SemSentSum* with three types of model relying on either 1) sentence or document embeddings 2) various hand-crafted features or 3) additional data.

Model	ROUGE-1	ROUGE-2
MMR	35.49	7.50
PV-DBOW+BS	36.10	6.77
PG-MMR	36.42	9.36
SVR	36.18	9.34
G-Flow	37.33	8.74
Peer 65	37.88	9.18
R2N2	38.16	9.52
TCSum	38.27	9.66
SemSentSum	39.12	9.59

Table 1: Comparison of various models using ROUGE-1/ROUGE-2 on DUC 2004 with 665 bytes summaries.

1. For the first category, we significantly outperform **MMR** (Bennani-Smires et al., 2018), **PV-DBOW+BS** (Mani et al., 2017) and **PG-MMR** (Lebanoff et al., 2018). Although their methods are based on embeddings to represent the meaning, it shows that using only various distance metrics or encoder-decoder architecture on these is not efficient for the task of multi-document summarization (as also shown in the Ablation Study). We hypothesize that *SemSentSum* performs better by leveraging pre-trained sentence embeddings and hence lowering the effects of data scarcity.
2. Systems based on hand-crafted features include a widely-used learning-based summarization method, built on support vector regression **SVR** (Li et al., 2007); a graph-based method based on approximating discourse graph **G-Flow** (Christensen et al., 2013); **Peer 65** which is the best peer systems participating in DUC evaluations; and the recursive neural network **R2N2** of Cao et al. (2015) that learns automatically combinations of hand-crafted features. As can be seen, among these models completely dependent on hand-crafted features, *SemSentSum* achieves highest performance on both ROUGE scores. This denotes that using different linguistic and word-based features might not be enough to capture the semantic structures, in addition to being cumbersome to craft.
3. The last type of model is shown in **TC-Sum** (Cao et al., 2017) and uses transfer learning from a text classifier model, based on a domain-related dataset of 30 000

Model	ROUGE-1	ROUGE-2
FreqSum	35.30	8.11
TsSum	35.88	8.15
Cont. LexRank	35.95	7.47
Centroid	36.41	7.97
CLASSY04	37.62	8.96
CLASSY11	37.22	9.20
GreedyKL	37.98	8.53
RegSum	38.57	9.75
GCN+PADG	38.23	9.48
SemSentSum	38.72	9.69

Table 2: Comparison of various models using ROUGE-1/2 on DUC 2004 with 100 words summaries.

documents from New York Times (sharing the same topics of the DUC datasets). In terms of ROUGE-1, *SemSentSum* significantly outperforms **TCSum** and performs similarly on ROUGE-2 score. This demonstrates that collecting more manually annotated data and training two models is unnecessary, in addition to being difficult to use in other domains, whereas *SemSentSum* is fully data driven, domain-independent and usable in realistic scenarios.

Table 2 depicts models producing 100 words summaries, all depending on hand-crafted features. We use as baselines **FreqSum** (Nenkova et al., 2006); **TsSum** (Conroy et al., 2006); traditional graph-based approaches such as **Cont. LexRank** (Erkan and Radev, 2004); **Centroid** (Radev et al., 2004); **CLASSY04** (Conroy et al., 2004); its improved version **CLASSY11** (Conroy et al., 2011) and the greedy model **GreedyKL** (Haghihi and Vanderwende, 2009). All of these models are significantly underperforming compared to *SemSentSum*. In addition, we include state-of-the-art models : **RegSum** (Hong and Nenkova, 2014) and **GCN+PADG** (Yasunaga et al., 2017). We outperform both in terms of ROUGE-1. For ROUGE-2 scores we achieve better results than **GCN+PADG** but without any use of domain-specific hand-crafted features and a much smaller and simpler model. Finally, **RegSum** achieves a similar ROUGE-2 score but computes sentence saliences based on word scores, incorporating a rich set of word-level and domain-specific features. Nonetheless, our model is competitive and does not depend on hand-crafted features due to its full data-driven nature and thus, it is not limi-

ted to a single domain.

Consequently, the experiments show that achieving good performance for multi-document summarization without hand-crafted features or additional data is clearly feasible and *SemSentSum* produces competitive results without depending on these, is domain independent, fast to train and thus usable in real scenarios.

Sentence Semantic Relation Graph Table 3 shows the results of different methods to create the sentence semantic relation graph with various thresholds t_{sim}^g for 665 bytes summaries (we obtain similar results for 100 words). A first observation is that using cosine similarity with sentence embeddings significantly outperforms all other methods for ROUGE-1 and ROUGE-2 scores, mainly because it relies on the semantic of sentences instead of their individual words. A second is that different methods evolve similarly : *PADG*, *Textrank*, *Tf-idf* behave similarly to an U-shaped curve for both ROUGE scores while *Cosine* is the only one having an inverted U-shaped curve. The reason for this behavior is a consequence of its distribution being similar to a normal distribution because it relies on the semantic instead of words, while the others are more skewed towards zero. This confirms our hypothesis that 1) having a complete graph does not allow the model to leverage much the semantic 2) a sparse graph might not contain enough information to exploit similarities. Finally, *Lexrank* and *ADG* have different trends between both ROUGE scores.

Ablation Study We quantify the contribution of each module of *SemSentSum* in Table 4 for 665 bytes summaries (we obtain similar results for 100 words). Removing the sentence encoder produces slightly lower results. This shows that the sentence semantic relation graph captures semantic attributes well, while the fine-tuned sentence embeddings obtained via the encoder help boost the performance, making these methods complementary. By disabling only the graph convolutional layer, a drastic drop in terms of performance is observed, which emphasizes that the relationship among sentences is indeed important and not present in the data itself. Therefore, our sentence semantic relation graph is able to capture sentence similarities by analyzing the semantic structures. Interestingly, if we remove the sentence encoder in addition to the graph convolutional layer, simi-

lar results are achieved. This confirms that alone, the sentence encoder is not able to compute an efficient representation of sentences for the task of multi-document summarization, probably due to the poor size of the DUC datasets. Finally, we can observe that the use of sentence embeddings only results in similar performance to the baselines, which rely on sentence or document embeddings (Bennani-Smires et al., 2018; Mani et al., 2017).

4 Related Work

The idea of using multiple embeddings has been employed at the word level. Kiela et al. (2018) use an attention mechanism to combine the embeddings for each word for the task of natural language inference. Xu et al. (2018); Bollegala et al. (2015) concatenate the embeddings of each word into a vector before feeding a neural network for the tasks of aspect extraction and sentiment analysis. To our knowledge, we are the first to combine multiple types of sentence embeddings.

Extractive multi-document summarization has been addressed by a large range of approaches. Several of them employ graph-based methods. Radev (2000) introduced a cross-document structure theory, as a basis for multi-document summarization. Erkan and Radev (2004) proposed LexRank, an unsupervised multi-document summarizer based on the concept of eigenvector centrality in a graph of sentences. Other works exploit shallow or deep features from the graph’s topology (Wan and Yang, 2006; Antiqueira et al., 2009). Wan and Yang (2008) pairs graph-based methods (e.g. random walk) with clustering. Mei et al. (2010) improved results by using a reinforced random walk model to rank sentences and keep non-redundant ones. The system by Christensen et al. (2013) does sentence selection, while balancing coherence and salience and by building a graph that approximates discourse relations across sentences (Mann and Thompson, 1988).

Besides graph-based methods, other viable approaches include Maximum Marginal Relevance (Carbonell and Goldstein, 1998), which uses a greedy approach to select sentences and considers the tradeoff between relevance and redundancy; support vector regression (Li et al., 2007); conditional random field (Galley, 2006); or hidden markov model (Conroy et al., 2004). Yet other approaches rely on n-grams regression as in Li et al.

Method	t_{sim}^g	ROUGE-1				t_{sim}^g	ROUGE-2			
		0.0	0.25	0.5	0.75		0.0	0.25	0.5	0.75
Cosine		38.49*	38.61*	39.12	35.54*		9.11*	9.07*	9.59	7.12*
Tf-idf		36.80*	36.23*	35.26*	35.71*		7.84*	7.78*	7.07*	7.46*
Textrank		35.66*	34.75*	35.41*	35.69*		7.83*	7.17*	7.20*	7.54*
Lexrank		37.04*	36.43*	36.27*	35.65*		7.90*	8.01*	7.64*	7.61*
ADG		35.48*	34.79*	34.78*	35.40*		6.96*	7.03*	7.01*	7.32*
PADG		36.81*	36.23*	35.26*	35.71*		7.84*	7.78*	7.07*	7.46*

Table 3: ROUGE-1/2 for various methods to build the sentence semantic relation graph. A score significantly different (according to a Welch Two Sample t-test, $p = 0.001$) than cosine similarity ($t_{sim}^g = 0.5$) is denoted by *.

Model	ROUGE-1	ROUGE-2
<i>SemSentSum</i>	39.12	9.59
- w/o Sent	38.38*	9.11*
- w/o GCN	35.88*	7.33*
- w/o GCN,Sent	35.89*	7.24*

Table 4: Ablation test. *Sent* is the sentence encoder and *GCN* the graph convolutional network. According to a Welch Two Sample t-test ($p = 0.001$), a score significantly different than *SemSentSum* is denoted by *.

(2013). More recently, Cao et al. (2015) built a recursive neural network, which tries to automatically detect combination of hand-crafted features. Cao et al. (2017) employ a neural model for text classification on a large manually annotated dataset and apply transfer learning for multi-document summarization afterward.

The work most closely related to ours is (Yasunaga et al., 2017). They create a normalized version of the approximate discourse graph (Christensen et al., 2013), based on hand-crafted features, where sentence nodes are normalized over all the incoming edges. They then employ a deep neural network, composed of a sentence encoder, three graph convolutional layers, one document encoder and an attention mechanism. Afterward, they greedily select sentences using TF-IDF similarity to detect redundant sentences. Our model differs in four ways : 1) we build our sentence semantic relation graph by using pre-trained sentence embeddings with cosine similarity, where neither heavy preprocessing, nor hand-crafted features are necessary. Thus, our model is fully data-driven and domain-independent unlike other systems. In addition, the sentence semantic relation graph could be used for other tasks than multi-document summarization, such as detecting information cascades, query-focused summarization, keyphrase extraction or information retrieval, as it

is not composed of hand-crafted features. 2) *SemSentSum* is much smaller and consequently has fewer parameters as it only uses a sentence encoder and a single convolutional layer. 3) The loss function is based on ROUGE-1 F_1 score instead of recall to prevent the tendency of choosing longer sentences. 4) Our method for summary generation is also different and novel as we leverage sentence embeddings to compute the similarity between a candidate sentence and the current summary instead of TF-IDF based approaches.

5 Conclusion

In this work, we propose a method to combine two types of sentence embeddings : 1) universal embeddings, pre-trained on a large corpus such as Wikipedia and incorporating general semantic structures across sentences and 2) domain-specific embeddings, learned during training. We merge them together by using a graph convolutional network that eliminates the need of hand-crafted features or additional annotated data.

We introduce a fully data-driven model *SemSentSum* that achieves competitive results for multi-document summarization on both kind of summary length (665 bytes and 100 words summaries), without requiring hand-crafted features or additional annotated data.

As *SemSentSum* is domain-independent, we believe that our sentence semantic relation graph and model can be used for other tasks including detecting information cascades, query-focused summarization, keyphrase extraction and information retrieval. In addition, we plan to leave the weights of the sentence semantic relation graph dynamic during training, and to integrate an attention mechanism directly into the graph.

Acknowledgments

We thank Michaela Benk for proofreading and helpful advice.

References

- Lucas Antiquiera, Osvaldo N. Oliveira, Luciano da Fontoura Costa, and Maria das Graças Volpe Nunes. 2009. [A complex network approach to text summarization](#). *Information Sciences*, 179(5) :584 – 599. Special Section - Quantum Structures : Theory and Applications.
- Kamil Bennani-Smires, Claudiu-Cristian Musat, Andreea Hossmann, Michael Baeriswyl, and Martin Jaggi. 2018. Simple unsupervised keyphrase extraction using sentence embeddings. *Proceedings of the 22nd Conference on Computational Natural Language Learning, CoNLL 2018, Brussels, Belgium, October 31 - November 1, 2018*, pages 221–229.
- Danushka Bollegala, Takanori Maehara, and Ken ichi Kawarabayashi. 2015. Unsupervised cross-domain word representation learning. *Proc. of the Annual Conference of the Association for Computational Linguistics (ACL) and International Joint Conferences on Natural Language Processing (IJCNLP)*.
- Ziqiang Cao, Wenjie Li, Sujian Li, and Furu Wei. 2017. [Improving multi-document summarization via text classification](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3053–3059.
- Ziqiang Cao, Furu Wei, Li Dong, Sujian Li, and Ming Zhou. 2015. [Ranking with recursive neural networks and its application to multi-document summarization](#). In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI’15*, pages 2153–2159. AAAI Press.
- Jaime Carbonell and Jade Goldstein. 1998. [The use of mmr, diversity-based reranking for reordering documents and producing summaries](#). In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’98*, pages 335–336, New York, NY, USA. ACM.
- Janara Christensen, Stephen Soderland, Oren Etzioni, et al. 2013. Towards coherent multi-document summarization. In *Proceedings of the 2013 conference of the North American chapter of the association for computational linguistics : Human language technologies*, pages 1163–1173.
- Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. 2016. [Fast and accurate deep network learning by exponential linear units \(elus\)](#). *International Conference on Learning Representations*.
- John M Conroy, Judith D Schlesinger, Jade Goldstein, and Dianne P O’leary. 2004. Left-brain/right-brain multi-document summarization. In *Proceedings of the Document Understanding Conference (DUC 2004)*.
- John M Conroy, Judith D Schlesinger, Jeff Kubina, Peter A Rankel, and Dianne P O’Leary. 2011. Classy 2011 at tac : Guided and multi-lingual summaries and evaluation metrics. *TAC*, 11 :1–8.
- John M. Conroy, Judith D. Schlesinger, and Dianne P. O’Leary. 2006. [Topic-focused multi-document summarization using an approximate oracle score](#). In *Proceedings of the COLING/ACL on Main Conference Poster Sessions, COLING-ACL ’06*, pages 152–159, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Günes Erkan and Dragomir R Radev. 2004. Lexrank : Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research*, 22 :457–479.
- Michel Galley. 2006. [A skip-chain conditional random field for ranking meeting utterances by importance](#). In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing, EMNLP ’06*, pages 364–372, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Aria Haghighi and Lucy Vanderwende. 2009. [Exploring content models for multi-document summarization](#). In *Proceedings of Human Language Technologies : The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, NAACL ’09*, pages 362–370, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Comput.*, 9(8) :1735–1780.
- Kai Hong and Ani Nenkova. 2014. Improving the estimation of word importance for news multi-document summarization. pages 712–721.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization : Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pages 448–456.
- Douwe Kiela, Changhan Wang, and Kyunghyun Cho. 2018. Dynamic meta-embeddings for improved sentence representations. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). *International Conference on Learning Representations*.
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#). *International Conference on Learning Representations*.
- Logan Lebanoff, Kaiqiang Song, and Fei Liu. 2018. Adapting the neural encoder-decoder framework from single to multi-document summarization. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Chen Li, Xian Qian, and Yang Liu. 2013. Using supervised bigram-based ilp for extractive summarization. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1 : Long Papers)*, volume 1, pages 1004–1013.

- Sujian Li, You Ouyang, Wei Wang, and Bin Sun. 2007. Multi-document summarization using support vector regression. In *In Proceedings of DUC*. Citeseer.
- C. Y. Lin. 2004. ROUGE : A package for automatic evaluation of summaries. In *Proceedings of the Workshop on Text Summarization Branches Out (WAS)*, Barcelona, Spain.
- Kaustubh Mani, Ishan Verma, and Lipika Dey. 2017. Multi-document summarization using distributed bag-of-words model. *arXiv preprint arXiv :1710.02745*.
- William C. Mann and Sandra A. Thompson. 1988. Rhetorical structure theory : Toward a functional theory of text organization. *Text*, 8(3) :243–281.
- Qiaozhu Mei, Jian Guo, and Dragomir Radev. 2010. Divrank: The interplay of prestige and diversity in information networks. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '10*, pages 1009–1018, New York, NY, USA. ACM.
- R. Mihalcea and P. Tarau. 2004. TextRank : Bringing order into texts. In *Proceedings of EMNLP-04 and the 2004 Conference on Empirical Methods in Natural Language Processing*.
- Ani Nenkova, Lucy Vanderwende, and Kathleen McKeown. 2006. 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, SIGIR '06*, pages 573–580, New York, NY, USA. ACM.
- Karolina Owczarzak, John M. Conroy, Hoa Trang Dang, and Ani Nenkova. 2012. An assessment of the accuracy of automatic evaluation in summarization. In *Proceedings of Workshop on Evaluation Metrics and System Comparison for Automatic Summarization*, pages 1–9, Stroudsburg, PA, USA. Association for Computational Linguistics.
- L. Page, S. Brin, R. Motwani, and T. Winograd. 1998. The pagerank citation ranking: Bringing order to the web. In *Proceedings of the 7th International World Wide Web Conference*, pages 161–172, Brisbane, Australia.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2018. Unsupervised Learning of Sentence Embeddings using Compositional n-Gram Features. In *NAACL 2018 - Conference of the North American Chapter of the Association for Computational Linguistics*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Dragomir R. Radev. 2000. A common theory of information fusion from multiple text sources step one: Cross-document structure. In *Proceedings of the 1st SIGDIAL Workshop on Discourse and Dialogue - Volume 10, SIGDIAL '00*, pages 74–83, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Dragomir R Radev, Hongyan Jing, Małgorzata Styś, and Daniel Tam. 2004. Centroid-based summarization of multiple documents. *Information Processing & Management*, 40(6) :919–938.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1) :1929–1958.
- Xiaojun Wan and Jianwu Yang. 2006. Improved affinity graph based multi-document summarization. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume : Short Papers, NAACL-Short '06*, pages 181–184, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Xiaojun Wan and Jianwu Yang. 2008. Multi-document summarization using cluster-based link analysis. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '08*, pages 299–306, New York, NY, USA. ACM.
- Hu Xu, Bing Liu, Lei Shu, and Philip S Yu. 2018. Double embeddings and cnn-based sequence labeling for aspect extraction. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Michihiro Yasunaga, Rui Zhang, Kshitijh Meelu, Ayush Pareek, Krishnan Srinivasan, and Dragomir R. Radev. 2017. Graph-based neural multi-document summarization. In *Proceedings of CoNLL-2017*. Association for Computational Linguistics.