

Reinforced Extractive Summarization with Question-Focused Rewards

Kristjan Arumae Fei Liu

Department of Computer Science

University of Central Florida, Orlando, FL 32816, USA

k.arumae@knights.ucf.edu feiliu@cs.ucf.edu

Abstract

We investigate a new training paradigm for extractive summarization. Traditionally, human abstracts are used to derive gold-standard labels for extraction units. However, the labels are often inaccurate, because human abstracts and source documents cannot be easily aligned at the word level. In this paper we convert human abstracts to a set of Cloze-style comprehension questions. System summaries are encouraged to preserve salient source content useful for answering questions and share common words with the abstracts. We use reinforcement learning to explore the space of possible extractive summaries and introduce a question-focused reward function to promote concise, fluent, and informative summaries. Our experiments show that the proposed method is effective. It surpasses state-of-the-art systems on the standard summarization dataset.

1 Introduction

We study extractive summarization in this work where salient word sequences are extracted from the source document and concatenated to form a summary (Nenkova and McKeown, 2011). Existing supervised approaches to extractive summarization frequently use human abstracts to create annotations for extraction units (Gillick and Favre, 2009; Li et al., 2013; Cheng and Lapata, 2016). E.g., a source word is labelled 1 if it appears in the abstract, 0 otherwise. Despite the usefulness, there are two issues with this scheme. First, a vast majority of the source words are tagged 0s, only a small portion are 1s. This is due to the fact that human abstracts are short and concise; they often contain words not present in the source. Second,

Source Document

The first doses of the Ebola vaccine were on a commercial flight to West Africa and were expected to arrive on Friday, according to a spokesperson from GlaxoSmithKline (GSK) one of the companies that has created the vaccine with the National Institutes of Health.

Another vaccine from Merck and NewLink will also be tested.

“Shipping the vaccine today is a major achievement and shows that we remain on track with the accelerated development of our candidate Ebola vaccine,” Dr. Moncef Slaoui, chairman of global vaccines at GSK said in a company release. (Rest omitted.)

Abstract

The first vials of an Ebola vaccine should land in Liberia Friday

Questions

Q: The first vials of an ____ vaccine should land in Liberia Friday

Q: The first vials of an Ebola vaccine should ____ in Liberia Friday

Q: The first vials of an Ebola vaccine should land in ____ Friday

Table 1: Example source document, the top sentence of the abstract, and system-generated Cloze-style questions. Source content related to the abstract is *italicized*.

not all labels are accurate. Source words that are labelled 0 may be paraphrases, generalizations, or otherwise related to words in the abstracts. These source words are often mislabelled. Consequently, leveraging human abstracts to provide supervision for extractive summarization remains a challenge.

Neural abstractive summarization can alleviate this issue by allowing the system to either copy words from the source texts or generate new words from a vocabulary (Rush et al., 2015; Nallapati et al., 2016; See et al., 2017). While the techniques are promising, they face other challenges, such as ensuring the summaries remain faithful to the original. Failing to reproduce factual details has been revealed as one of the main obstacles for neural abstractive summarization (Cao et al., 2018; Song and Liu, 2018). This study thus chooses to focus on neural extractive summarization.

We explore a new training paradigm for extractive summarization. We convert human abstracts to a set of Cloze-style comprehension questions, where the question body is a sentence of the abstract with a blank, and the answer is an entity or a keyword. Table 1 shows an example. Because the

questions cannot be answered by applying general world knowledge, system summaries are encouraged to preserve salient source content that is relevant to the questions (\approx human abstract) such that the summaries can work as a document surrogate to predict correct answers. We use an attention mechanism to locate segments of a summary that are relevant to a given question so that the summary can be used to answer multiple questions.

This study extends the work of (Lei et al., 2016) to use reinforcement learning to explore the space of extractive summaries. While the original work focuses on generating rationales to support supervised classification, the goal of our study is to produce fluent, generic document summaries. The question-answering (QA) task is designed to fulfill this goal and the QA performance is only secondary. Our research contributions can be summarized as follows:

- we investigate an alternative training scheme for extractive summarization where the summaries are encouraged to be semantically close to human abstracts in addition to sharing common words;
- we compare two methods to convert human abstracts to Cloze-style questions and investigate its impact on QA and summarization performance. Our results surpass those of previous systems on a standard summarization dataset.

2 Related Work

This study focuses on generic summarization. It is different from the query-based summarization (Daumé III and Marcu, 2006; Dang and Owczarzak, 2008), where systems are trained to select text pieces related to *predefined* queries. In this work we have no predefined queries but the system carefully generates questions from human abstracts and learns to produce generic summaries that are capable of answering all questions.

Cloze questions have been used in reading comprehension (Richardson et al., 2013; Weston et al., 2016; Mostafazadeh et al., 2016; Rajpurkar et al., 2016) to test the system’s ability to perform reasoning and language understanding. Hermann et al. (2015) describe an approach to extract (context, question, answer) triples from news articles. Our work draws on this approach to automatically create questions from human abstracts.

Reinforcement learning (RL) has been recently applied to a number of NLP applications, includ-

ing dialog generation (Li et al., 2017), machine translation (MT) (Ranzato et al., 2016; Gu et al., 2018), question answering (Choi et al., 2017), and summarization and sentence simplification (Zhang and Lapata, 2017; Paulus et al., 2017; Chen and Bansal, 2018; Narayan et al., 2018). This study leverages RL to explore the space of possible extractive summaries. The summaries are encouraged to preserve salient source content useful for answering questions as well as sharing common words with the abstracts.

3 Our Approach

Given a source document X , our system generates a summary $Y = (y_1, y_2, \dots, y_{|Y|})$ by identifying consecutive sequences of words: y_t is 1 if the t -th source word is included in the summary, 0 otherwise. In this section we investigate a question-oriented reward $\mathcal{R}(Y)$ that encourages summaries to contain sufficient content useful for answering key questions about the document (§3.1); we then use reinforcement learning to explore the space of possible extractive summaries (§3.2).

3.1 Question-Focused Reward

We reward a summary if it can be used as a document surrogate to answer important questions. Let $\{(Q_k, e_k^*)\}_{k=1}^K$ be a set of question-answer pairs for a source document, where e_k^* is the ground-truth answer corresponding to an entity or a keyword. We encode the question Q_k into a vector: $\mathbf{q}_k = \text{Bi-LSTM}(Q_k) \in \mathbb{R}^d$ using a bidirectional LSTM, where the last outputs of the forward and backward passes are concatenated to form a question vector. We use the same Bi-LSTM to encode the summary Y to a sequence of vectors: $(\mathbf{h}_1^S, \mathbf{h}_2^S, \dots, \mathbf{h}_{|S|}^S) = \text{Bi-LSTM}(Y)$, where $|S|$ is the number of words in the summary; $\mathbf{h}_t^S \in \mathbb{R}^d$ is the concatenation of forward and backward hidden states at time step t . Figure 1 provides an illustration of the system framework.

An attention mechanism is used to locate parts of the summary that are relevant to Q_k . We define $\alpha_{k,i} \propto \exp(\mathbf{q}_k \mathbf{W}^a \mathbf{h}_i^S)$ to represent the importance of the i -th summary word (\mathbf{h}_i^S) to answering the k -th question (\mathbf{q}_k), characterized by a bilinear term (Chen et al., 2016a). A context vector \mathbf{c}_k is constructed as a weighted sum of all summary words relevant to the k -th question, and it is used to predict the answer. We define the QA reward $\mathcal{R}_a(Y)$ as the log-likelihood of correctly predict-

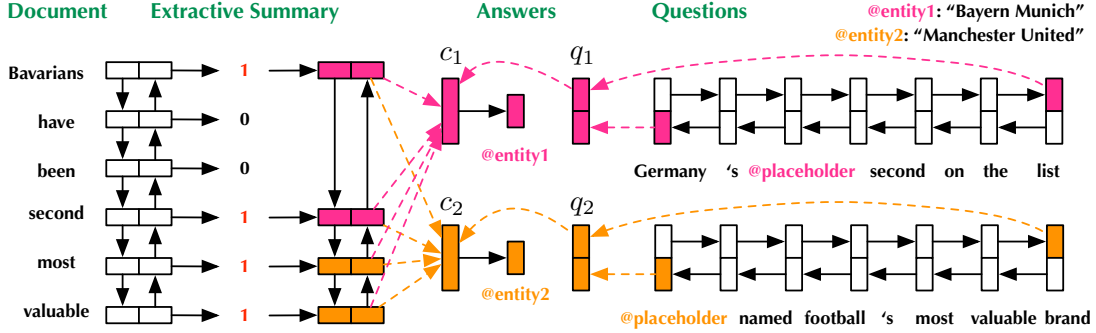


Figure 1: System framework. The model uses an extractive summary as a document surrogate to answer important questions about the document. The questions are automatically derived from the human abstract.

ing all answers. $\{\mathbf{W}^a, \mathbf{W}^c\}$ are learnable model parameters.

$$\alpha_{k,i} = \frac{\exp(\mathbf{q}_k \mathbf{W}^a \mathbf{h}_i^S)}{\sum_{i=1}^{|S|} \exp(\mathbf{q}_k \mathbf{W}^a \mathbf{h}_i^S)} \quad (1)$$

$$\mathbf{c}_k = \sum_{i=1}^{|S|} \alpha_{k,i} \mathbf{h}_i^S \quad (2)$$

$$P(e_k|Y, Q_k) = \text{softmax}(\mathbf{W}^c \mathbf{c}_k) \quad (3)$$

$$\mathcal{R}_a(Y) = \frac{1}{K} \sum_{k=1}^K \log P(e_k^*|Y, Q_k) \quad (4)$$

In the following we describe approaches to obtain a set of question-answer pairs $\{(Q_k, e_k^*)\}_{k=1}^K$ from a human abstract. In fact, this formulation has the potential to make use of multiple human abstracts (subject to availability) in a unified framework; in that case, the QA pairs will be extracted from all abstracts. According to Eq. (4), the system is optimized to generate summaries that preserve salient source content sufficient to answer *all questions* (\approx human abstract).

We expect to harvest one question-answer pair from each sentence of the abstract. More are possible, but the QA pairs will contain duplicate content. There are a few other noteworthy issues. If we do not collect any QA pairs from a sentence of the abstract, its content will be left out of the system summary. It is thus crucial for the system to extract at least one QA pair from *any* sentence in an automatic manner. Further, the questions must not be answered by simply applying general world knowledge. We expect the adequacy of the summary to have a direct influence on whether or not the questions will be correctly answered. Motivated by these considerations, we perform the following steps. We split a human abstract to a set of sentences, identify an answer token from each

sentence, then convert the sentence to a question by replacing the token with a placeholder, yielding a Cloze question. We explore two approaches to extract answer tokens:

- *Entities*. We extract four types of named entities $\{\text{PER}, \text{LOC}, \text{ORG}, \text{MISC}\}$ from sentences and treat them as possible answer tokens.
- *Keywords*. This approach identifies the ROOT word of a sentence dependency parse tree and treats it as a keyword-based answer token. Not all sentences contain entities, but every sentence has a root word; it is often the main verb of the sentence.

We obtain K question-answer pairs from each human abstract, one pair per sentence. If there are less than K sentences in the abstract, the QA pairs of the top sentences will be duplicated, with the assumption that the top sentences are more important than others. If multiple entities reside in a sentence, we randomly pick one as the answer token; otherwise if there are no entities, we use the root word instead.

To ensure that the extractive summaries are concise, fluent, and close to the original wording, we add additional components to the reward function: (i) we define $\mathcal{R}_s(Y) = \frac{1}{|Y|} \sum_{t=1}^{|Y|} y_t - \delta$ to restrict the summary size. We require the percentage of selected source words to be close to a predefined threshold δ . This constraint works well at restricting length, with the average summary size adhering to this percentage; (ii) we further introduce $\mathcal{R}_f(Y) = \sum_{t=2}^{|Y|} |y_t - y_{t-1}|$ to encourage the summaries to be fluent. This component is adopted from (Lei et al., 2016), where few 0/1 switches between y_{t-1} and y_t indicates the system is selecting consecutive word sequences; (iii) we encourage system and reference summaries to share common bigrams. This practice has shown suc-

cess in earlier studies (Gillick and Favre, 2009). $\mathcal{R}_b(Y)$ is defined as the percentage of reference bigrams successfully covered by the system summary. These three components together ensure the well-formedness of extractive summaries. The final reward function $\mathcal{R}(Y)$ is a linear interpolation of all the components; γ, α, β are coefficients and we describe their parameter tuning in §4.

$$\mathcal{R}(Y) = \mathcal{R}_a(Y) + \gamma \mathcal{R}_b(Y) - \alpha \mathcal{R}_f(Y) - \beta \mathcal{R}_s(Y) \quad (5)$$

3.2 Reinforcement Learning

In the following we seek to optimize a policy $P(Y|X)$ for generating extractive summaries so that the expected reward $\mathbb{E}_{P(Y|X)}[\mathcal{R}(Y)]$ is maximized. Taking derivatives of this objective with respect to model parameters θ involves repeatedly sampling summaries $\hat{Y} = (\hat{y}_1, \hat{y}_2, \dots, \hat{y}_{|Y|})$ (illustrated in Eq. (6)). In this way reinforcement learning exploits the space of extractive summaries of a source document.

$$\begin{aligned} & \nabla_{\theta} \mathbb{E}_{P(Y|X)}[\mathcal{R}(Y)] \\ &= \mathbb{E}_{P(Y|X)}[\mathcal{R}(Y) \nabla_{\theta} \log P(Y|X)] \\ &\approx \frac{1}{N} \sum_{n=1}^N \mathcal{R}(\hat{Y}^{(n)}) \nabla_{\theta} \log P(\hat{Y}^{(n)}|X) \quad (6) \end{aligned}$$

To calculate $P(Y|X)$ and then sample \hat{Y} from it, we use a bidirectional LSTM to encode a source document to a sequence of vectors: $(\mathbf{h}_1^D, \mathbf{h}_2^D, \dots, \mathbf{h}_{|X|}^D) = \text{Bi-LSTM}(X)$. Whether to include the t -th source word in the summary (\hat{y}_t) thus can be decided based on \mathbf{h}_t^D . However, we also want to accommodate the previous $t-1$ sampling decisions ($\hat{y}_{1:t-1}$) to improve the fluency of the extractive summary. Following (Lei et al., 2016), we introduce a single-direction LSTM encoder whose hidden state \mathbf{s}_t tracks the sampling decisions up to time step t (Eq. 8). It represents the semantic meaning encoded in the current summary. To sample the t -th word, we concatenate the two vectors $[\mathbf{h}_t^D || \mathbf{s}_{t-1}]$ and use it as input to a feed-forward layer with sigmoid activation to estimate $\hat{y}_t \sim P(y_t | \hat{y}_{1:t-1}, X)$ (Eq. 7).

$$P(y_t | \hat{y}_{1:t-1}, X) = \sigma(\mathbf{W}^h [\mathbf{h}_t^D || \mathbf{s}_{t-1}] + b^h) \quad (7)$$

$$\mathbf{s}_t = \text{LSTM}([\mathbf{h}_t^D || \hat{y}_t], \mathbf{s}_{t-1}) \quad (8)$$

$$P(\hat{Y}|X) = \prod_{t=1}^{|Y|} P(\hat{y}_t | \hat{y}_{1:t-1}, X) \quad (9)$$

Note that Eq. (7) can be pretrained using goldstandard summary sequence $Y^* = (y_1^*, y_2^*, \dots, y_{|Y|}^*)$ to minimize the word-level cross-entropy loss,

System	R-1	R-2	R-L
LSA (Steinberger and Jezek, 2004)	21.2	6.2	14.0
LexRank (Erkan and Radev, 2004)	26.1	9.6	17.7
TextRank (Mihalcea and Tarau, 2004)	23.3	7.7	15.8
SumBasic (Vanderwende et al., 2007)	22.9	5.5	14.8
KL-Sum (Haghighi and Vanderwende, 2009)	20.7	5.9	13.7
Distraction-M3 (Chen et al., 2016b)	27.1	8.2	18.7
Seq2Seq w/ Attn (See et al., 2017)	25.0	7.7	18.8
Pointer-Gen w/ Cov (See et al., 2017)	29.9	10.9	21.1
Graph-based Attn (Tan et al., 2017)	30.3	9.8	20.0
Extr+EntityQ (this paper)	31.4	11.5	21.7
Extr+KeywordQ (this paper)	31.7	11.6	21.5

Table 2: Results on the CNN test set (full-length F1 scores).

where we set y_t^* as 1 if (x_t, x_{t+1}) is a bigram in the human abstract. For reinforcement learning, our goal is to optimize the policy $P(Y|X)$ using the reward function $\mathcal{R}(Y)$ (§3.1) during the training process. Once the policy $P(Y|X)$ is learned, we do not need the reward function (or any QA pairs) at test time to generate generic summaries. Instead we choose \hat{y}_t that yields the highest probability $\hat{y}_t = \arg \max P(y_t | \hat{y}_{1:t-1}, X)$.

4 Experiments

All training, validation, and testing was performed using the CNN dataset (Hermann et al., 2015; Nallapati et al., 2016) containing news articles paired with human-written highlights (i.e., abstracts). We observe that a source article contains 29.8 sentences and an abstract contains 3.54 sentences on average. The train/valid/test splits contain 90,266, 1,220, 1,093 articles respectively.

4.1 Hyperparameters

The hyperparameters, tuned on the validation set, include the following: the hidden state size of the Bi-LSTM is 256; the hidden state size of the single-direction LSTM encoder is 30. Dropout rate (Srivastava, 2013), used twice in the sampling component, is set to 0.2. The minibatch size is set to 256. We apply early stopping on the validation set, where the maximum number of epochs is set to 50. Our source vocabulary contains 150K words; words not in the vocabulary are replaced by the $\langle \text{unk} \rangle$ token. We use 100-dimensional word embeddings, initialized by GloVe (Pennington et al., 2014) and remain trainable. We set $\beta = 2\alpha$ and select the best $\alpha \in \{\underline{10}, 20, 50\}$ and $\gamma \in \{5, 6, 7, \underline{8}\}$ using the valid set (best value underlined). The maximum length of input is set to 100 words; δ is set to be 0.4 (≈ 40 words). We use the Adam optimizer (Kingma and Ba, 2015) with

an initial learning rate of $1e-4$ and halve the learning rate if the objective worsens beyond a threshold ($> 10\%$). As mentioned we utilized a bigram based pretraining method. We found that this stabilized the training of the full model.

4.2 Results

We compare our methods with state-of-the-art published systems, including both extractive and abstractive approaches (their details are summarized below). We experiment with two variants of our approach. “EntityQ” uses QA pairs whose answers are named entities. “KeywordQ” uses pairs whose answers are sentence root words. According to the R-1, R-2, and R-L scores (Lin, 2004) presented in Table 2, both methods are superior to the baseline systems on the benchmark dataset, yielding 11.5 and 11.6 R-2 F-scores, respectively.

- **LSA** (Steinberger and Jezek, 2004) uses the latent semantic analysis technique to identify semantically important sentences.
- **LexRank** (Erkan and Radev, 2004) is a graph-based approach that computes sentence importance based on the concept of eigenvector centrality in a graph representation of source sentences.
- **TextRank** (Mihalcea and Tarau, 2004) is an unsupervised graph-based ranking algorithm inspired by algorithms PageRank and HITS.
- **SumBasic** (Vanderwende et al., 2007) is an extractive approach that assumes words occurring frequently in a document cluster have a higher chance of being included in the summary.
- **KL-Sum** (Haghighi and Vanderwende, 2009) describes a method that greedily adds sentences to the summary so long as it decreases the KL divergence.
- **Distraction-M3** (Chen et al., 2016b) trains the summarization model to not only attend to specific regions of input documents, but also distract the attention to traverse different content of the source document.
- **Pointer-Generator** (See et al., 2017) allows the system to not only copy words from the source text via pointing but also generate novel words through the generator.
- **Graph-based Attention** (Tan et al., 2017) introduces a graph-based attention mechanism to enhance the encoder-decoder framework.

	K1	K2	K3	K4	K5
# Uniq Entities	23.7K	37.0K	46.1K	50.3K	50.3K
Train Acc (%)	46.1	37.2	34.2	33.6	34.8
Valid Acc (%)	12.8	14.0	14.7	15.7	15.4
Valid R-2 F (%)	11.2	11.1	11.2	11.1	10.8
# Uniq Keywds	7.3K	10.4K	12.5K	13.7K	13.7K
Train Acc (%)	30.5	28.2	27.6	27.5	27.5
Valid Acc (%)	19.3	22.5	22.2	23.0	21.9
Valid R-2 F (%)	11.2	11.1	10.8	11.0	10.8

Table 3: Train/valid accuracy and R-2 F-scores when using varying numbers of QA pairs ($K=1$ to 5) in the reward func.

In Table 3, we vary the number of QA pairs used per article in the reward function ($K=1$ to 5). The summaries are encouraged to contain comprehensive content useful for answering *all* questions. When more QA pairs are used ($K1 \rightarrow K5$), we observe that the number of answer tokens has increased and almost doubled: 23.7K ($K1$) \rightarrow 50.3K ($K5$) for entities as answers, and 7.3K \rightarrow 13.7K for keywords. The enlarged answer space has an impact on QA accuracies. When using entities as answers, the training accuracy is 34.8% ($Q5$) and validation is 15.4% ($Q5$), and there appears to be a considerable gap between the two. In contrast, the gap is quite small when using keywords as answers (27.5% and 21.9% for $Q5$), suggesting that using sentence root words as answers is a more viable strategy to create QA pairs.

Comparing to QA studies (Chen et al., 2016a), we remove the constraint that requires answer entities (or keywords) to reside in the source documents. Adding this constraint improves the QA accuracy for a standard QA system. However, because our system does not perform QA during testing (the question-answer pairs are not available for the test set) but only generate generic summaries, we do not enforce this requirement and report no testing accuracies. We observe that the R-2 scores only present minor changes from $K1$ to $K5$. We conjecture that more question-answer pairs do not make the summaries contain more comprehensive content because the input and the summary are relatively short; $K=1$ yields the best results.

In Table 4, we present example system and reference summaries. Our extractive summaries can be overlaid with the source documents to assist people with browsing through the documents. In this way the summaries stay true to the original and do not contain information that was not in the source documents.

Future work. We are interested in investigating

Source Document
It was all set for a fairytale ending for record breaking jockey AP McCoy . In the end it was a different but familiar name who won the Grand National on Saturday . 25-1 outsider Many Clouds, who had shown little form going into the race, won by a length and a half, ridden by jockey Leighton Aspell . Aspell won last year's Grand National too, making him the first jockey since the 1950s to ride back-to-back winners on different horses . "It feels wonderful, I asked big questions," Aspell said...
Abstract
25-1 shot Many Clouds wins Grand National Second win a row for jockey Leighton Aspell First jockey to win two in a row on different horses since 1950s

Table 4: Example system summary and human abstract. The summary words are shown in bold in the source document.

approaches that automatically group selected summary segments into clusters. Each cluster can capture a unique aspect of the document, and clusters of text segments can be color-highlighted. Inspired by the recent work of Narayan et al. (2018), we are also interested in conducting the usability study to test how well the summary highlights can help users quickly answer key questions about the documents. This will provide an alternative strategy for evaluating our proposed method against both extractive and abstractive baselines.

5 Conclusion

In this paper we explore a new training paradigm for extractive summarization. Our system converts human abstracts to a set of question-answer pairs. We use reinforcement learning to exploit the space of extractive summaries and promote summaries that are concise, fluent, and adequate for answering questions. Results show that our approach is effective, surpassing state-of-the-art systems.

Acknowledgments

We thank the anonymous reviewers for their valuable suggestions. This work is in part supported by an unrestricted gift from Bosch Research. Kristjan Arumae gratefully acknowledges a travel grant provided by the National Science Foundation.

References

Ziqiang Cao, Furu Wei, Wenjie Li, and Sujian Li. 2018. Faithful to the original: Fact aware neural abstractive summarization. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.

Danqi Chen, Jason Bolton, and Christopher D. Manning. 2016a. A thorough examination of the cnn/daily mail reading comprehension task. In *Pro-*

ceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL).

- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, and Hui Jiang. 2016b. Distraction-based neural networks for document summarization. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)*.
- Yen-Chun Chen and Mohit Bansal. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Jianpeng Cheng and Mirella Lapata. 2016. Neural summarization by extracting sentences and words. In *Proceedings of ACL*.
- Eunsol Choi, Daniel Hewlett, Jakob Uszkoreit, Illia Polosukhin, Alexandre Lacoste, and Jonathan Berant. 2017. Coarse-to-fine question answering for long documents. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Hoa Trang Dang and Karolina Owczarzak. 2008. Overview of the TAC 2008 update summarization task. In *Proceedings of Text Analysis Conference*.
- Hal Daumé III and Daniel Marcu. 2006. Bayesian query-focused summarization. In *Proceedings of the 44th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Günes Erkan and Dragomir R. Radev. 2004. LexRank: Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research*.
- Dan Gillick and Benoit Favre. 2009. A scalable global model for summarization. In *Proceedings of the NAACL Workshop on Integer Linear Programming for Natural Language Processing*.
- Jiatao Gu, Daniel Jiwoong Im, and Victor O.K. Li. 2018. Neural machine translation with gumbel-greedy decoding. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI)*.
- Aria Haghighi and Lucy Vanderwende. 2009. Exploring content models for multi-document summarization. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proceedings of Neural Information Processing Systems (NIPS)*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Chen Li, Fei Liu, Fuliang Weng, and Yang Liu. 2013. Document summarization via guided sentence compression. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Jiwei Li, Will Monroe, Tianlin Shi, Sebastien Jean, Alan Ritter, and Dan Jurafsky. 2017. Adversarial learning for neural dialogue generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Chin-Yew Lin. 2004. ROUGE: a package for automatic evaluation of summaries. In *Proceedings of ACL Workshop on Text Summarization Branches Out*.
- Rada Mihalcea and Paul Tarau. 2004. TextRank: Bringing order into text. In *Proceedings of EMNLP*.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Caglar Gulcehre, and Bing Xiang. 2016. Abstractive text summarization using sequence-to-sequence RNNs and beyond. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning (CoNLL)*.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Ranking sentences for extractive summarization with reinforcement learning. In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Ani Nenkova and Kathleen McKeown. 2011. Automatic summarization. *Foundations and Trends in Information Retrieval*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the Conference Empirical Methods in Natural Language Processing (EMNLP)*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence level training with recurrent neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Matthew Richardson, Christopher J.C. Burges, and Erin Renshaw. 2013. MCTest: A challenge dataset for the open-domain machine comprehension of text. In *Proceedings of Empirical Methods in Natural Language Processing (EMNLP)*.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for sentence summarization. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Kaiqiang Song and Fei Liu. 2018. Structure-infused copy mechanisms for abstractive summarization. In *Proceedings of the International Conference on Computational Linguistics (COLING)*.
- Nitish Srivastava. 2013. *Improving Neural Networks with Dropout*. Master’s thesis, University of Toronto, Toronto, Canada.
- Josef Steinberger and Karel Jezek. 2004. Using latent semantic analysis in text summarization and summary evaluation. In *Proceedings of ISIM*.
- Jiwei Tan, Xiaojun Wan, and Jianguo Xiao. 2017. Abstractive document summarization with a graph-based attentional neural model. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Lucy Vanderwende, Hisami Suzuki, Chris Brockett, and Ani Nenkova. 2007. Beyond SumBasic: Task-focused summarization with sentence simplification and lexical expansion. *Information Processing and Management* 43(6):1606–1618.
- Jason Weston, Antoine Bordes, Sumit Chopra, Sasha Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2016. Towards AI-complete question answering: A set of prerequisite toy tasks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Xingxing Zhang and Mirella Lapata. 2017. Sentence simplification with deep reinforcement learning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.