

# Data-driven Summarization of Scientific Articles

Nikola I. Nikolov, Michael Pfeiffer, Richard H.R. Hahnloser

Institute of Neuroinformatics, University of Zürich and ETH Zürich, Switzerland

{niniko, pfeiffer, rich}@ini.ethz.ch

## Abstract

Data-driven approaches to sequence-to-sequence modelling have been successfully applied to short text summarization of news articles. Such models are typically trained on input-summary pairs consisting of only a single or a few sentences, partially due to limited availability of multi-sentence training data. Here, we propose to use scientific articles as a new milestone for text summarization: large-scale training data come almost for free with two types of high-quality summaries at different levels - the title and the abstract. We generate two novel multi-sentence summarization datasets from scientific articles and test the suitability of a wide range of existing extractive and abstractive neural network-based summarization approaches. Our analysis demonstrates that scientific papers are suitable for data-driven text summarization. Our results could serve as valuable benchmarks for scaling sequence-to-sequence models to very long sequences.

## 1. Introduction

The goal of automatic text summarization is to produce a shorter, informative version of an input text. While extractive summarization only consists of selecting important sentences from the input, abstractive summarization *generates* content without explicitly re-using whole sentences (Nenkova et al., 2011). Text summarization is an area with much promise in today’s age of information overflow. In the domain of scientific literature, the rate of publications grows exponentially (Hunter and Cohen, 2006), which calls for efficient automatic summarization tools.

Recent state-of-the-art summarization methods learn to summarize in a data-driven way, relying on large collections of input-summary training examples. The majority of previous work focused on short summarization of news articles, such as to generate a title (Rush et al., 2015; Nallapati et al., 2016). One major challenge is to scale these methods to process long input/output sequence pairs. Currently, availability of large-scale high-quality training data is scarce.

In this paper, we explore the suitability of **scientific journal articles** as a new benchmark for data-driven text summarization. The typical well-structured format of scientific papers makes them an interesting challenge, and provides plenty of freely available training data, because every article comes with a summary in the form of its abstract, and, in even more compressed form, its title. We make a first step towards summarization of whole scientific articles, by composing two novel large datasets for **scientific summarization**: title-abstract pairs (*title-gen*), composed of 5 million papers in the biomedical domain, and abstract-body pairs (*abstract-gen*) composed of 900k papers<sup>1</sup>. The second dataset is particularly challenging, because it is intended for summarizing the full body of the paper in terms of the abstract (the lengths of input/output sequences are substantially longer than what has been considered so far in previous research, see Table 1).

We evaluate a range of existing state-of-the-art approaches on these datasets: extractive approaches based on word embeddings, as well as word, subword, and character-level

encoder-decoder models that use recurrent as well as convolutional modules. We perform a quantitative and qualitative analysis of the models’ outputs.

## 2. Background

### 2.1. Extractive Summarization

Given an input document consisting of  $T_s$  sentences  $\mathbf{s} = \{s_1, \dots, s_{T_s}\}$ , the goal of **extractive summarization** is to select the  $K$  most salient sentences as the output summary. Extractive summarization typically involves a **sentence representation module**  $e$ , that represents each input sentence  $s_i$  in a common space as  $r_i = e(s_i)$ , e.g. as a vector of real numbers; as well as a **ranking module**  $score$ , that weights the salience  $w_i = score(r_i)$  of each sentence. A typical approach to unsupervised extractive summarization is to implement  $w_i$  as the similarity between  $r_i$  and a document representation (or a **document centroid**)  $r_d = e(d)$  (Radev et al., 2004). Alternatively, one can compute  $w_i$  as the **sentence centrality**, which is an adjacency-based measure of sentence importance (Erkan and Radev, 2004).

In this work, we propose two simple unsupervised baselines for extractive summarization, both of which rely on word embeddings (Mikolov et al., 2013). The first, *tfidf-emb*, represents each sentence in the input document as the weighted sum of its constituent word embeddings, similar to (Rossiello et al., 2017):

$$r_i = e(s_i) = \frac{1}{t_i} \sum_{x \in s_i} t(x) \cdot E(x), \quad (1)$$

where  $E(x)$  is the embedding of word  $x$ ,  $t(x)$  is an (optional) weighting function that weighs the importance of a word, and  $t_i = \sum_{x \in s_i} t(x)$  is a normalization factor. As a weighing function, we use the term-frequency inverse document frequency (TF-IDF) score, similar to (Brokos et al., 2016). Each sentence embedding  $r_i$  can then be ranked by computing its cosine similarity  $sim(r_d, r_i)$  to a document centroid  $r_d$ , computed similarly as  $r_i$ . The summary consists of the top  $K$  sentences with embeddings most similar to the document embedding.

The second baseline, *rwmd-rank*, ranks the salience of a sentence in terms of its similarity to all the other sentences

<sup>1</sup>Both datasets are available at <https://github.com/ninikolov/data-driven-summarization>, including versions with and without preprocessing.

in the document. All similarities are stored in an intra-sentence similarity matrix  $W$ . We use the Relaxed Word Mover’s Distance (RWMD) to compute this matrix (Kusner et al., 2015):

$$W_{ij} = \text{rwmd}(s_i, s_j) = \max(\text{rwmd}_p(s_i, s_j), \text{rwmd}_p(s_j, s_i))$$

$$\text{rwmd}_p(s_i, s_j) = \sum_{x \in s_i} \min_{x' \in s_j} \text{dist}(E(x), E(x')), \quad (2)$$

where  $s_i$  and  $s_j$  are two sentences and  $\text{dist}(E(x), E(x'))$  is the Euclidean distance between the embeddings of words  $x$  and  $x'$  in the sentences. To rank the sentences, we apply the graph-based method from the LexRank system (Erkan and Radev, 2004). LexRank represents the input as a highly connected graph, in which vertices represent sentences, and edges between sentences are assigned weights equal to their similarity from  $W$ . The centrality of a sentence is then computed using the PageRank algorithm (Page et al., 1999).

## 2.2. Abstractive Summarization

Given an input sequence of  $T_x$  words  $\mathbf{x} = \{x_1, \dots, x_{T_x}\}$  coming from a fixed-length input vocabulary  $V_x$  of size  $K_x$ , the goal of abstractive summarization is to produce a condensed sequence of  $T_y$  summary words  $\mathbf{y} = \{y_1, \dots, y_{T_y}\}$  from a summarization vocabulary  $V_y$  of size  $K_y$ , where  $T_x \gg T_y$ . Abstractive summarization is a structured prediction problem that can be solved by learning a probabilistic mapping  $p(\mathbf{y}|\mathbf{x}, \theta)$  for the summary  $\mathbf{y}$ , given the input sequence  $\mathbf{x}$  (Dietterich et al., 2008):

$$p(\mathbf{y}|\mathbf{x}, \theta) = \prod_i^{T_y} p(y_i | \{y_0, \dots, y_{i-1}\}, \mathbf{x}, \theta). \quad (3)$$

The **encoder-decoder architecture** is a recently proposed general framework for structured prediction (Cho et al., 2015), in which the distribution  $\arg \max_{\mathbf{y}} p(\mathbf{y}|\mathbf{x}, \theta)$  is learned using two neural networks: an **encoder** network  $e$ , which produces intermediate representations of the input, and a **decoder** language modelling network  $d$ , which generates the target summary. The decoder is conditioned on a context vector  $c$ , which is recomputed from the encoded representation at each decoding step. The encoder-decoder was first implemented using Recurrent Neural Networks (RNNs) (Sutskever et al., 2014; Cho et al., 2014) that process the input sequentially. Recent studies have shown that convolutional neural networks (CNNs) (LeCun et al., 1998) can outperform RNNs in sequence transduction tasks (Kalchbrenner et al., 2016; Gehring et al., 2017). Unlike RNNs, CNNs can be efficiently implemented on parallel GPU hardware. This advantage is particularly important when working with very long input and output sequences, such as whole paragraphs or documents. CNNs create hierarchical representations over the input in which lower layers operate on nearby elements and higher layers implement increasing levels of abstraction.

In this work, we investigate the performance of three existing systems that operate on different levels of sequence

granularity. The first, *lstm*, is a recurrent Long Short Term Memory (LSTM) encoder-decoder model (Sutskever et al., 2014) with an attention mechanism (Bahdanau et al., 2014) that operates on the word level, processing the input sequentially. The second system, *fconv*, is a convolutional encoder-decoder model from (Gehring et al., 2017). *fconv* works on the subword level and segments words into smaller units using the byte pair encoding scheme. Using subword units improves the generation quality when dealing with rare or unknown words (Sennrich et al., 2015). The third system, *c2c*, is a character-level encoder-decoder model from (Lee et al., 2016) that models  $\mathbf{x}$  and  $\mathbf{y}$  as individual characters, with no explicit segmentation between tokens. *c2c* first builds representations of groups of characters in the input using a series of convolutional layers. It then applies a recurrent encoder-decoder, similar to the *lstm* system.

## 2.3. Scientific Articles

Previous research on summarization of scientific articles has focused almost exclusively on extractive methods (Nenkova et al., 2011). In (Lloret et al., 2013), the authors develop an unsupervised system for abstract generation of biomedical papers, that first selects relevant content from the body, following which it performs an abstractive information fusion step. More recently, (Kim et al., 2016) consider the problem of supervised generation of sentence-level summaries for each paragraph of the introduction of a paper. They construct a training dataset of computer science papers from arXiv, selecting the most informative sentence as the summary of each paragraph using the Jaccard similarity. Thus, their target summary is fully contained in the input. In (Collins et al., 2017), they develop a supervised extractive summarization framework which they apply to a dataset of 10k computer science papers. To the best of our knowledge, our work is the first on abstractive title generation of scientific articles, and is the first to consider supervised generation of the abstract directly from the full body of the paper. The datasets we utilize here are also substantially larger than in previous work on scientific summarization.

Scientific articles are potentially more challenging to summarize than news articles because of their compact, implicit discourse style (Biber and Gray, 2010). While the events described by news headlines frequently recur in related articles, a scientific title focuses on the unique contribution that sets a paper apart from previous research (Teufel and Moens, 2002). Furthermore, while the first two sentences of a news article are typically sufficiently informative to generate its headline (Nallapati et al., 2016; Teufel and Moens, 2002), the first sentences of the abstract or introduction of a paper typically contain background information on the research topic. Constructing a good scientific title thus requires understanding and integrating concepts from multiple sentences of the abstract.

## 3. Datasets

To investigate the performance of encoder-decoder neural networks as generative models of scientific text, we constructed two novel datasets for scientific summarization.

Table 1: Statistics (mean and standard deviation) of the two scientific summarization datasets: *title-gen* and *abstract-gen*. Token/sentence counts are computed with NLTK.

<i>title-gen</i>	<i>Abstract</i>	<i>Title</i>
<b>Token count</b>	245 ± 54	15 ± 4
<b>Sentence count</b>	14 ± 4	1
<b>Sent. token count</b>	26 ± 14	-
<b>Overlap</b>	73% ± 18%	
<b>Repeat</b>	44% ± 11%	-
<b>Size (tr/val/test)</b>	5'000'000/6844/6935	

<i>abstract-gen</i>	<i>Body</i>	<i>Abstract</i>
<b>Token count</b>	4600 ± 1987	254 ± 54
<b>Sentence count</b>	172 ± 78	10 ± 3
<b>Sent. token count</b>	26 ± 17	26 ± 14
<b>Overlap</b>	68% ± 10%	
<b>Repeat</b>	74% ± 7%	44% ± 11%
<b>Size (tr/val/test)</b>	893'835/10'916/10'812	

For *title-gen* we used MEDLINE<sup>2</sup>, whereas for *abstract-gen* we used the PubMed open access subset<sup>3</sup>. MEDLINE contains scientific metadata in XML format of  $\sim 25$  million papers in the biomedical domain, whereas the PubMed open access subset contains metadata and full text of  $\sim 1.5$  million papers.

We processed the XML files<sup>4</sup> to pair the abstract of a paper to its title (*title-gen* dataset) or the full body (*abstract-gen*), skipping any figures, tables or section headings in the body. We then apply several preprocessing steps from the MOSES statistical machine translation pipeline<sup>5</sup>, including tokenization and conversion to lowercase. Any URLs were removed, all numbers replaced with #, and any pairs with abstract lengths not in the range of 150-370 tokens, title lengths not within 6-25 tokens, and body lengths not within 700-10000 tokens were excluded.

The *Overlap*  $o(\mathbf{x}, \mathbf{y}) = \frac{|\{\mathbf{y}\} \cap \{\mathbf{x}\}|}{|\{\mathbf{y}\}|}$  is the fraction of unique output (summary) tokens  $\mathbf{y}$  that overlap with an input token  $\mathbf{x}$  (excluding punctuation and stop words). As can be seen in Table 1, the overlaps are large in our datasets, indicating frequent reuse of words. The *Repeat*  $e(\mathbf{s}) = \frac{\sum_i o(\bar{\mathbf{s}}_i, \mathbf{s}_i)}{|\mathbf{s}|}$  is the average overlap of each sentence  $\mathbf{s}_i$  in a text with the remainder of the text (where  $\bar{\mathbf{s}}_i$  denotes the complement of sentence  $\mathbf{s}_i$ ). *Repeat* measures the redundancy of content within a text: a high value indicates frequent repetition of content. Whereas in abstracts there are only moderate levels of repetition, in the bodies the repetition rates are much higher, possibly because concepts and ideas are reiterated in multiple sections of the paper.

## 4. Evaluation Set-up

We evaluated the performance of several state-of-the-art approaches on our scientific summarization datasets. The extractive systems we consider are: *lead*, *lexrank*, *tfidf-emb*, and *rwmd-rank*. The *lead* baseline returns the first sentence of the abstract for *title-gen*, or the first 10 sentences of the body for *abstract-gen*. *lexrank* (Erkan and Radev, 2004) is a graph-based centrality approach frequently used as a baseline in the literature. *emb-tfidf* uses sentence embeddings<sup>6</sup> to select the most salient sentences from the input,

while *rwmd-rank* uses the Relaxed Word Mover’s Distance (as described in Section 2.1.). *oracle* estimates an upper bound for the extractive summarization task by finding the most similar sentence in the input document for each sentence in the original summary. We use the Relaxed Word Mover’s Distance to compute the output of the oracle.

The abstractive systems we consider are: *lstm*, *fconv*, and *c2c*, described in Section 2.2.. For *lstm*, we set the input/output vocabularies to 80000, use two LSTM layers of 1000 hidden units each, and word embedding of dimension 500 (we found no improvement from additionally increasing the size of this model). For *c2c* and *fconv*, we use the default hyper-parameters that come with the public implementations provided by the authors of the systems. The *title-gen lstm*, *c2c*, and *fconv* were trained for 11, 8, and 20 epochs, respectively, until convergence.

We were unable to train *lstm* and *c2c* on *abstract-gen* because of the very high memory and time requirements associated with the recurrent layers in these models. We found *fconv* to be much more efficient to train, and we succeeded in training a default model for 17 epochs. For *title-gen*, we used beam search with beam size 20, while for *abstract-gen* we found a beam size of 5 to perform better.

### 4.1. Quantitative Evaluation

In Tables 2 and 3, we evaluate our approaches using the ROUGE metric (Lin, 2004), which is a recall-based metric frequently used for summarization, and METEOR (Denkowski and Lavie, 2014), which is a precision-based metric for machine translation. Overlap can be interpreted as the tendency of the model to directly copy input content instead of generating novel correct or incorrect words; whereas Repeat measures a model’s tendency to repeat itself, which is a frequent issue with encoder-decoder models (Suzuki and Nagata, 2017).

On title generation, *rwmd-rank* achieved the best performance in terms of selecting a sentence as the title. In overall, the abstractive systems significantly outperformed the extractive systems, as well as the extractive oracle. *c2c* and *fconv* performed much better than *lstm*, with a very high rate of overlap. The ROUGE performance of *c2c* and *fconv* is similar, despite the difference of a few R-2 points in favour of *fconv* (that model is evaluated on a subword-level ground truth file, where we observe a slight increase of 1-2 ROUGE points on average due to the conversion).

On abstract generation, the *lead-10* baseline remained tough to beat in terms of ROUGE, and only the extractive systems managed to surpass it by a small margin. All extractive systems achieved similar results, with *rwmd-rank*

<sup>2</sup>[https://nlm.nih.gov/databases/download/pubmed\\_medline.html](https://nlm.nih.gov/databases/download/pubmed_medline.html)

<sup>3</sup><https://ncbi.nlm.nih.gov/pmc/tools/openftlist>

<sup>4</sup>We use [https://titipata.github.io/pubmed\\_parser](https://titipata.github.io/pubmed_parser).

<sup>5</sup>[github.com/moses-smt/mosesdecoder](https://github.com/moses-smt/mosesdecoder)

<sup>6</sup>We use the best-performing Word2Vec model from (Chiu et al., 2016), which is trained on PubMed and MEDLINE.

Table 2: Metric results for the *title-gen* dataset. R-1, R-2, R-L represent the ROUGE-1/2/L metrics.

Model	R-1	R-2	R-L	METEOR	Overlap	Token count
<i>oracle</i>	0.386	0.184	0.308	0.146	-	29 ± 14
<i>lead-1</i>	0.218	0.061	0.169	0.077	-	28 ± 14
<i>lexrank</i>	0.26	0.089	0.201	0.089	-	32 ± 14
<i>emb-tfidf</i>	0.252	0.081	0.193	0.082	-	35 ± 17
<i>rwmd-rank</i>	0.311	0.13	0.245	0.116	-	28 ± 13
<i>lstm</i>	0.375	0.173	0.329	0.204	78% ± 20%	12 ± 3
<i>c2c</i>	<b>0.479</b>	0.264	<b>0.418</b>	0.237	93% ± 10%	14 ± 4
<i>fconv</i>	0.463	<b>0.277</b>	0.412	<b>0.27</b>	95% ± 9%	15 ± 7

Table 3: Metric results for the *abstract-gen* dataset. R-1, R-2, R-L represent the ROUGE-1/2/L metrics.

Model	R-1	R-2	R-L	METEOR	Overlap	Repeat	Token count
<i>oracle</i>	0.558	0.266	0.316	0.214	-	42% ± 10%	327 ± 99
<i>lead-10</i>	0.385	0.111	0.18	0.138	-	20% ± 4%	312 ± 88
<i>lexrank</i>	0.45	<b>0.163</b>	0.213	0.157	-	52% ± 10%	404 ± 131
<i>emb-tfidf</i>	0.445	0.159	0.216	0.159	-	52% ± 10%	369 ± 117
<i>rwmd-rank</i>	<b>0.454</b>	0.159	<b>0.216</b>	0.167	-	50% ± 10%	344 ± 93
<i>fconv</i>	0.354	0.131	0.209	<b>0.212</b>	98% ± 2%	52% ± 28%	194 ± 15

having a minor edge, while the abstractive *fconv* performed poorly, even though it performed best in terms of METEOR. We observed a much higher repeat rate in the output summaries than the observed 44% average in the original abstracts (Table 1). As revealed by the large Repeat standard deviation for *fconv*, some examples are affected by very frequent repetitions.

## 4.2. Qualitative Evaluation

In Tables 4 and 5, we present two shortened inputs from our *title-gen* and *abstract-gen* test sets, along with original and system-generated summaries. In Figure 1, we show a histogram of the locations of input sentences, that estimates which locations were most preferred on average when producing a summary.

We observe a large variation in the sentence locations selected by the extractive systems on *title-gen* (Figure 1a), with the first sentence having high importance. Based on our inspection, it is rare that a sentence from the abstract will match the title exactly - the title is also typically shorter than an average sentence from the abstract (Table 1). A good title seems to require the selection, combination and paraphrasing of suitable parts from multiple sentences, as also shown by the original titles in our examples. Many of the titles generated by the abstractive systems sound faithful, and at first glance can pass for a title of a scientific paper. The abstractive models are good at discerning important from unimportant content in the abstract, at extracting long phrases, or sometimes whole sentences, and at abstractively combining the information to generate a title. *lstm* is more prone to generate novel words, whereas *c2c* and *fconv* mostly rely on direct copying of content from the abstract, as also indicated by their overlap scores.

Closer inspection of the titles reveals occasional subtle mistakes: for example, in the first example in Table 4, the *fconv* model incorrectly selected "scopolamine- and cisplatin-induced" which was investigated in the previous work of the authors and is not the main focus of the article. The model also copied the incorrect genus, "mouse" instead of "rat".

Sometimes the generated titles sound too general, and fail to communicate the specifics of the paper: in the second example, all models produced "a model of basal ganglia", missing to include the keyword "reinforcement learning": "a model of reinforcement learning in the basal ganglia". These mistakes highlight the complexity of the task, and show that there is still much room for further improvement.

As shown in Figure 1b, the introductory and concluding sections are often highly relevant for abstract generation, however relevant content is spread across the entire paper. Interestingly, in the example in Table 5, there is a wide range of content that was selected by the extractive systems, with little overlap across systems. For instance, *rwmd-rank* overlaps with *oracle* by 3 sentences, and only by 1 sentence with *emb-tfidf*. The outputs of the abstractive *fconv* system on abstract generation are poor in quality, and many of the generated abstracts lack coherent structure and content flow. There is also frequent repetition of entire sentences, as shown by the last sentences produced by *fconv* in Table 5. *fconv* also appears to only use the first 60 sentences of the paper to construct the abstract (Figure 1c).

## 5. Conclusion

We evaluated a range of extractive and abstractive neural network-based summarization approaches on two novel datasets constructed from scientific journal articles. While the results for title generation are promising, the models struggled with generating the abstract. This difficulty highlights the necessity for developing novel models capable of efficiently dealing with long input and output sequences, while at the same time preserving the quality of generated sentences. We hope that our datasets will promote more work in this area. A direction to explore in future work is hybrid extractive-abstractive end-to-end approaches that jointly select content and then paraphrase it to produce a summary.

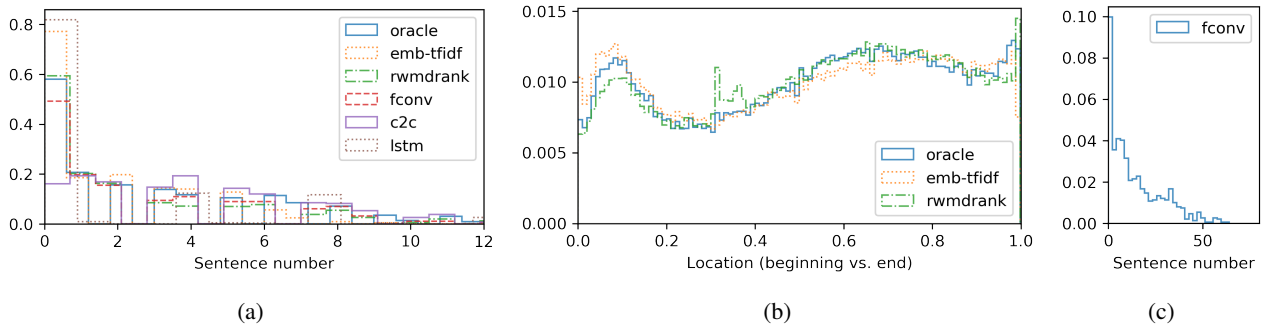


Figure 1: Sentence selection (normalized) histograms computed on the test set, showing the input locations that were most preferred on average by the systems on *title-gen* (a) and *abstract-gen* (b), (c). For (b), we normalize the sentence locations by the length of each paper, to get a better uniform view (there is a large variation in the length of a paper, as shown in Table 1). For the abstractive systems, we search for the closest sentences in the input using relaxed word mover’s distance (see Section 2.1.).

Table 4: Examples from the test set of *title-gen*. The outputs of the extractive systems are highlighted as: **oracle**, **tfidf-emb**, **rwm-rank**. For the abstractive systems, we manually highlighted the text of the concepts that are relevant for the task (errors are highlighted in red).

<p><i>Example 1 (Giridharan et al., 2015) Abstract:</i> <b>Amyloid (A)-induced neurotoxicity is a major pathological mechanism of Alzheimers disease (AD).</b> <b>Our previous studies have demonstrated that schisandrin B (Sch B), an antioxidant lignan from Schisandra chinensis, could protect mouse brain against scopolamine- and cisplatin-induced neuronal dysfunction.</b> In the present study, we examined the <b>protective effect of Sch B</b> against intracerebroventricular (ICV)-infused A-induced neuronal dysfunction in rat cortex and explored the potential mechanism of its action. <b>Our results showed that 26 days co-administration of Sch B significantly improved the behavioral performance of A (140)-infused rats in step-through test.</b> At the same time, Sch B attenuated A-induced increases in oxidative and nitrosative stresses (...) The aforementioned effects of <b>Sch B</b> suggest its <b>protective role</b> against <b>A-induced neurotoxicity</b> through intervention in the negative cycle of RAGE-mediated A accumulation during AD patho-physiology.</p> <p><i>Original title:</i> <b>schisandrin b ameliorates icv-infused amyloid induced oxidative stress and neuronal dysfunction through inhibiting rage / nf-b / mapk and up-regulating hsp / beclin expression</b></p> <p><i>lstm:</i> <b>schisandrin b</b> , an antioxidant lignan from schisandra chinensis , <b>protects</b> against <b>amyloid -induced neurotoxicity</b></p> <p><i>c2c:</i> <b>schisandra chinensis b protects</b> against <b>intracerebroventricular-infused amyloid induced neuronal dysfunction in rat cortex</b></p> <p><i>fconv:</i> <b>schisandrin b protects mouse brain against scopolamine- and cisplatin- induced neurotoxicity in rats</b></p>
<p><i>Example 2 (Fee, 2012) Abstract:</i> <b>In its simplest formulation, reinforcement learning is based on the idea that if an action taken in a particular context is followed by a favorable outcome, then, in the same context, the tendency to produce that action should be strengthened, or reinforced.</b> (...) Recent experiments in the <b>songbird</b> suggest that vocal-related BG circuitry receives two functionally distinct excitatory inputs. (...) <b>The other is an efference copy of motor commands from a separate cortical brain region that generates vocal variability during learning.</b> Based on these findings, I propose here a <b>general model</b> of vertebrate BG function that combines context information with a distinct motor efference copy signal. (...) <b>The model makes testable predictions about the anatomical and functional properties of hypothesized context and efference copy inputs to the striatum from both thalamic and cortical sources.</b></p> <p><i>Original title:</i> oculomotor learning revisited : a <b>model of reinforcement learning in the basal ganglia</b> incorporating an efference copy of motor actions .</p> <p><i>lstm:</i> a <b>model</b> of basal ganglia function .</p> <p><i>c2c:</i> a <b>general model</b> of <b>vertebrate basal ganglia function</b> .</p> <p><i>fconv:</i> a <b>model</b> of basal ganglia function in the <b>songbird</b> .</p>

## 6. Acknowledgments

We thank the reviewers for their useful comments, and NVIDIA for the donation of a TITAN X graphics card.

## 7. References

- Bahdanau, D., Cho, K., and Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Biber, D. and Gray, B. (2010). Challenging stereotypes about academic writing: Complexity, elaboration, explicitness. *Journal of English for Academic Purposes*, 9(1):2–20.
- Brokos, G.-I., Malakasiotis, P., and Androutsopoulos, I. (2016). Using centroids of word embeddings and word mover’s distance for biomedical document retrieval in question answering. *arXiv preprint arXiv:1608.03905*.
- Chiu, B., Crichton, G., Korhonen, A., and Pyysalo, S. (2016). How to train good word embeddings for biomedical nlp. *Proceedings of BioNLP16*, page 166.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Cho, K., Courville, A., and Bengio, Y. (2015). Describing multimedia content using attention-based encoder-decoder networks. *arXiv preprint arXiv:1507.01053*.
- Collins, E., Augenstein, I., and Riedel, S. (2017). A supervised approach to extractive summarisation of scientific

Table 5: Two examples from the test set of *abstract-gen*. The outputs of the extractive systems are highlighted as: **tfidf-emb** and **rwmd-rank**, whereas **gray** denotes overlap between the two. In **bold** we mark the content that was selected by the *fconv* system (next page in full), and in underline we mark the selection of the *oracle*.

*Example 1 (Pyysalo et al., 2011) Body:* **In recent years, there has been a significant shift in focus in biomedical information extraction from simple pairwise relations representing associations such as protein-protein interactions (PPI) toward representations that capture typed, structured associations of arbitrary numbers of entities in specific roles, frequently termed event extraction [1].** Much of this work draws on the GENIA Event corpus (...) This resource served also as the source for the annotations in the first collaborative evaluation of biomedical event extraction methods, the 2009 BioNLP shared task on event extraction (BioNLP ST) [6] as well as for the GENIA subtask of the second task in the series [7, 8]. Another recent trend in the domain is a move toward the application of extraction methods to the full scale of the existing literature, with results for various targets covering the entire PubMed literature database of nearly 20 million citations being made available [9, 10, 11, 12]. **As event extraction methods** initially developed to target the set of events defined in the GENIA / BioNLP ST corpora are now being applied at PubMed scale, it makes sense to ask how much of the full spectrum of gene/protein associations found there they can maximally cover (...) By contrast, we will assume that associations not appearing in this data cannot be extracted: as the overwhelming majority of current event extraction methods **are based on supervised machine learning or hand-crafted rules written with reference to the annotated data**, it reasonable to assume as a first approximation that their coverage of associations not appearing in that data is zero. **In this study, we seek to characterize the full range of associations of specific genes/proteins described in the literature and estimate what coverage of these associations event extraction systems relying on currently available resources can maximally achieve.** To address these questions, **it is necessary not only to have an inventory of concepts that (largely) covers the ways in which genes/proteins can be associated, but also to be able to estimate the relative frequency with which these concepts are used to express gene/protein associations** in the literature. (...) Here, as we are interested in particular in texts describing associations between two or more gene/protein related entities, **we apply a focused selection, picking only those individual sentences in which** two or more mentions co-occur. **While this excludes associations in which the entities occur in different sentences, their relative frequency is expected to be low: for example, in the BioNLP ST data, all event participants occurred within a single sentence in 95% of the targeted biomolecular event statements.** (...) **Here, we follow the assumption that when two entities are stated to be associated in some way, the most important words expressing their association will typically be found on the shortest dependency path connecting the two entities (cf. the shortest path hypothesis of Bunescu and Mooney [30]).** The specific dependency representation (...) Table 3 shows the words most frequently occurring on these paths. **This list again suggests an increased focus on words relating to gene/protein associations: expression is the most frequent word on the paths, and binding appears in the top-ranked words.** (...) Finally, to make this pair data consistent with the TPS event spans, tokenization and other features, we aligned the entity annotations of the two corpora. (...) **This processing was applied to the BioNLP ST training set, creating a corpus of 6889 entity pairs of which 1119 (16%) were marked as expressing an association (positive)** (...) Evaluation. We first evaluated each of the word rankings discussed in the section on Identification of Gene/Protein Associations by comparing the ranked lists of words against the set of single words marked as trigger expressions in the BioNLP ST development data, (...) To evaluate the capability of the presented approach to identify new expressions of gene/protein associations, we next performed a manual study of candidate words for stating gene/protein associations using the E w ranking. (...) **We then selected the words ranked highest by E w that were not known, grouped by normalized and lemmatized form, and added for reference examples of frequent shortest dependency paths on which any of these words appear (see example in Table 5).** (...) If static relations and experimental observations and manipulations are excluded as (arguably) not in scope for event extraction, this estimate suggests that currently available resources for event extraction cover over 90% of all events involving gene/protein entities in PubMed. Discussion. **We found that out of all gene/protein associations in PubMed, currently existing resources for event extraction are lacking in coverage of a number of event types such as dissociation, many relatively rare (though biologically important) protein post-translational modifications, as well as some high-level process types involving genes/proteins such as apoptosis.** (...) This suggests that for practical applications it may be important to consider also this class of associations. (...) **While these results are highly encouraging, it must be noted that the approach to identifying gene/protein associations considered here is limited in a number of ways: it excludes associations stated across sentence boundaries and ones for which the shortest path hypothesis does not hold, does not treat multi-word expressions as wholes, ignores ambiguity in implicitly assuming a single sense for each word, and only directly includes associations stated between exactly two entities.** The approach is also fundamentally limited to associations expressed through specific words and thus blind to e.g. part-of relations implied by statements such as CD14 Sp1-binding site. (...) Conclusions. **We have presented an approach to discovering expressions of gene/protein associations from PubMed based on named entity co-occurrences, shortest dependency paths and an unlexicalized classifier to identify likely statements of gene/protein associations.** Drawing on the automatically created full-PubMed annotations of the Turku PubMed-Scale (TPS) corpus and using the BioNLP09 shared task data to define positive and negative examples of association statements, we distilled an initial set of over 30 million protein mentions into a set of 46,000 unique unlexicalized paths estimated likely to express gene/protein associations. These paths were then used to rank all words in PubMed by the expected number of times they are predicted to express such associations, and 1200 candidate association-expressing words not appearing in the BioNLP09 shared task data evaluated manually. **Study of these candidates suggested 18 new event classes for the GENIA ontology and indicated that the majority of statements of gene/protein associations not covered by currently available resources are not statements of biomolecular events but rather statements of static relations or experimental manipulation.** (...) **It could thus be assumed that the event types and the specific statements annotated in GENIA would have only modest coverage of all gene/protein association types and statements in PubMed.** (...)



*Example 1 Original abstract:* Background: Event extraction following the GENIA Event corpus and BioNLP shared task models has been a considerable focus of recent work in biomedical information extraction. This work includes efforts applying event extraction methods to the entire PubMed literature database, far beyond the narrow subdomains of biomedicine for which annotated resources for extraction method development are available. Results: In the present study, our aim is to estimate the coverage of all statements of gene/protein associations in PubMed that existing resources for event extraction can provide. We base our analysis on a recently released corpus automatically annotated for gene/protein entities and syntactic analyses covering the entire PubMed, and use named entity co-occurrence, shortest dependency paths and an unlexicalized classifier to identify likely statements of gene/protein associations. A set of high-frequency/high-likelihood association statements are then manually analyzed with reference to the GENIA ontology. Conclusions: We present a first estimate of the overall coverage of gene/protein associations provided by existing resources for event extraction. Our results suggest that for event-type associations this coverage may be over 90%. We also identify several biologically significant associations of genes and proteins that are not addressed by these resources, suggesting directions for further extension of extraction coverage.

*Example 1 fconv:* background in recent years , there has been a significant shift in focus in biomedical information extraction from simple pairwise relations representing associations such as protein-protein interactions ( ppi ) toward representations that capture typed , structured associations of arbitrary numbers of entities in specific roles , frequently termed event extraction . as event extraction methods are based on supervised machine learning or hand-crafted rules written with reference to the annotated data , it is necessary not only to have an inventory of concepts that ( largely ) covers the ways in which genes / proteins can be associated , but also to be able to estimate the relative frequency with which these concepts are used to express gene / protein associations . results we apply a focused selection , picking king only those individual sentences in which at least one named entity has been tagged . we apply a focused selection , picking only those individual sentences in which at least one named entity has been tagged . we apply a focused selection , picking only those individual sentences in which at least one named entity has been tagged . we apply a probabilistic approach

*Example 2 (Kumar et al., 2013) Body:* **Both alopecia areata (AA) and vitiligo are autoimmune diseases, and their coexistence in the same patient is not uncommon, as vitiligo has been reported to occur in 4.1 % of patients of AA and is about 4 times more common in patients with AA than in the general population.** [1] However, their colocalization over the same site is exceedingly rare, with less than five cases being reported in the literature.[2,3,4] **We present a case of a 15-year-old male child who had vitiligo and later developed AA over the existing lesions of vitiligo over face and scalp** and have attempted to elucidate the current understanding of mechanisms of coexistence of these two diseases. **A 12-year-old boy presented to the skin outpatient department with history of depigmented areas on the scalp, face, neck, arms and legs for 5 years. He also gave a history of development of patchy loss of hair over some of these lesions for 3 years.** There was no previous history of any trauma or medications. Family history was not relevant. **On examination, there were depigmented macules over the scalp, forehead, eyebrows, eyebrows, perioral, preauricular regions, neck, elbows, hands, feet, shins, nose, chin, hands, knees and feet. Patches of hair loss were seen, limited to some of these depigmented areas over the vertex and occipital region of the scalp and eyebrows [Figure 3]. Other body areas were not affected by patchy hair loss. Clinically, the diagnosis of vitiligo with AA was made. (..) Additionally, the basal layer of the epidermis was almost devoid of pigment, [Figure 5] confirming the diagnosis of vitiligo over the same site. (..) Both AA and vitiligo are clubbed under the spectrum of autoimmune disorders. (..) Our case lends support to the hypothesis that AA and vitiligo share a common pathogenic pathway including autoimmune response against some common antigens like those derived from the bulb melanocytes. Melanocytes-derived peptide antigens released during vitiligo pathogenesis could act as auto-antigens not only for vitiligo, but also for AA, and autoimmune Th-cells against them could also trigger a response against the hair follicle melanocytes, thus, pre-disposing to AA. (..text omitted..)**

*Example 2 Original abstract:* Both alopecia areata (AA) and vitiligo are autoimmune diseases, and their coexistence in the same patient is not uncommon, as vitiligo has been reported to occur in 4.1% of patients of AA. We present a case of a 15-year-old male child who had vitiligo and later developed AA over the existing lesions of vitiligo over face and scalp and have attempted to elucidate the current understanding of mechanisms of coexistence of these two diseases. Our case lends support to the hypothesis that AA and vitiligo share a common pathogenic pathway including autoimmune response against some common antigens like those derived from the bulb melanocytes. Stimulation of proinflammatory T-helper(Th)-1 cell mediated immunological response or inactivation of a suppressor T-cell mediated response could be the common underlying phenomenon. However, the striking rarity of colocalization of these two diseases has led to the recent debate over current understanding of their pathogenesis, and whether this association is merely a coincidence. As both AA and vitiligo are frequent and chronic dermatological disorders, it is of utmost importance to gain more understanding into their pathogenic mechanisms so that more definitive treatment modalities may be devised and the quality of life of these patients can be improved.

*Example 2 fconv:* alopecia areata ( aa ) and vitiligo are autoimmune diseases , and their coexistence in the same patient is not uncommon , as vitiligo has been reported to occur in 4.1 % of patients of aa and is about 4 times more common in patients with aa than in the general population . we present a case of a 15- year-old male child who had vitiligo and later developed aa over the scalp , forehead , eyebrows , eyebrows , perioral , preauricular regions , nose and chin depigmented macules over the scalp , forehead , eyebrows , eyebrows , eyebrows , perioral , preauricular regions , nose and chin depigmented macules over the scalp , forehead , eyebrows , periorbital , perioral , preauricular regions , nose and chin depigmented macules over the scalp , forehead , eyebrows , periorbital , perioral , preauricular regions , nose and chin depigmented macules over the scalp , forehead , eyebrows , periorbital

- papers. *arXiv preprint arXiv:1706.03946*.
- Denkowski, M. and Lavie, A. (2014). Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the EACL 2014 Workshop on Statistical Machine Translation*.
- Dietterich, T. G., Domingos, P., Getoor, L., Muggleton, S., and Tadepalli, P. (2008). Structured machine learning: the next ten years. *Machine Learning*, 73(1):3–23.
- Erkan, G. and Radev, D. R. (2004). Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research*, 22:457–479.
- Fee, M. S. (2012). Oculomotor learning revisited: a model of reinforcement learning in the basal ganglia incorporating an efference copy of motor actions. *Frontiers in neural circuits*, 6:38.
- Gehring, J., Auli, M., Grangier, D., Yarats, D., and Dauphin, Y. N. (2017). Convolutional sequence to sequence learning. *arXiv preprint arXiv:1705.03122*.
- Giridharan, V. V., Thandavarayan, R. A., Arumugam, S., Mizuno, M., Nawa, H., Suzuki, K., Ko, K. M., Krishnamurthy, P., Watanabe, K., and Konishi, T. (2015). Schisandrin b ameliorates icv-infused amyloid  $\beta$  induced oxidative stress and neuronal dysfunction through inhibiting rage/nf- $\kappa$ b/mapk and up-regulating hsp/beclin expression. *PLoS One*, 10(11):e0142483.
- Hunter, L. and Cohen, K. B. (2006). Biomedical language processing: what’s beyond pubmed? *Molecular cell*, 21(5):589–594.
- Kalchbrenner, N., Espeholt, L., Simonyan, K., van den Oord, A., Graves, A., and Kavukcuoglu, K. (2016). Neural machine translation in linear time. *CoRR*, abs/1610.10099.
- Kim, M., Singh, M. D., and Lee, M. (2016). Towards abstraction from extraction: Multiple timescale gated recurrent unit for summarization. *arXiv preprint arXiv:1607.00718*.
- Kumar, S., Mittal, J., and Mahajan, B. (2013). Colocalization of vitiligo and alopecia areata: coincidence or consequence? *International journal of trichology*, 5(1):50.
- Kusner, M. J., Sun, Y., Kolkin, N. I., Weinberger, K. Q., et al. (2015). From word embeddings to document distances. In *ICML*, volume 15, pages 957–966.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov.
- Lee, J., Cho, K., and Hofmann, T. (2016). Fully character-level neural machine translation without explicit segmentation. *arXiv preprint arXiv:1610.03017*.
- Lin, C.-Y. (2004). Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the ACL-04 workshop*, volume 8.
- Lloret, E., Romá-Ferri, M. T., and Palomar, M. (2013). Compendium: A text summarization system for generating abstracts of research papers. *Data & Knowledge Engineering*, 88:164–175.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Nallapati, R., Zhou, B., Gulçehre, Ç., and Xiang, B. (2016). Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*.
- Nenkova, A., Maskey, S., and Liu, Y. (2011). Automatic summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts of ACL 2011*, page 3. Association for Computational Linguistics.
- Page, L., Brin, S., Motwani, R., and Winograd, T. (1999). The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.
- Pyysalo, S., Ohta, T., and Tsujii, J. (2011). An analysis of gene/protein associations at pubmed scale. *Journal of biomedical semantics*, 2(5):S5.
- Radev, D. R., Jing, H., Styś, M., and Tam, D. (2004). Centroid-based summarization of multiple documents. *Information Processing & Management*, 40(6):919–938.
- Rossello, G., Basile, P., and Semeraro, G. (2017). Centroid-based text summarization through compositionality of word embeddings. *MultiLing 2017*, page 12.
- Rush, A. M., Chopra, S., and Weston, J. (2015). A neural attention model for abstractive sentence summarization. *arXiv preprint arXiv:1509.00685*.
- Sennrich, R., Haddow, B., and Birch, A. (2015). Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Suzuki, J. and Nagata, M. (2017). Cutting-off redundant repeating generations for neural abstractive summarization. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, volume 2, pages 291–297.
- Teufel, S. and Moens, M. (2002). Summarizing scientific articles: experiments with relevance and rhetorical status. *Computational linguistics*, 28(4):409–445.