

# Ranking sentences from product description & bullets for better search

Prateek Verma  
Jet.com and Walmart Labs  
Hoboken, New Jersey  
prateek.verma@jet.com

Aliasgar Kutiyawala  
Jet.com and Walmart Labs  
Hoboken, New Jersey  
aliasgar@jet.com

Ke Shen  
Jet.com and Walmart Labs  
Hoboken, New Jersey  
ke.shen@jet.com

## ABSTRACT

Products in an ecommerce catalog contain information-rich fields like description and bullets that can be useful to extract entities (attributes) using NER based systems. However, these fields are often verbose and contain lot of information that is not relevant from a search perspective. Treating each sentence within these fields equally can lead to poor full text match and introduce problems in extracting attributes to develop ontologies, semantic search etc. To address this issue, we describe two methods based on extractive summarization with reinforcement learning by leveraging information in product titles and search click through logs to rank sentences from bullets, description, etc. Finally, we compare the precision of these two models.

## CCS CONCEPTS

• **Information systems** → **Information extraction; Summarization.**

## KEYWORDS

Search and Ranking, Information Retrieval, Extractive Summarization, Reinforcement Learning, E-Commerce, Information Extraction

### ACM Reference Format:

Prateek Verma, Aliasgar Kutiyawala, and Ke Shen. 2019. Ranking sentences from product description & bullets for better search. In *Proceedings of ACM SIGIR Workshop on eCommerce (SIGIR 2019 eCom)*. ACM, New York, NY, USA, Article 4, 6 pages.

## 1 INTRODUCTION

Many search engine frameworks like Solr [7] and Elasticsearch [6] treat each sentence within a field in the document equally and this can lead to irrelevant documents present in the recall set. Consider Figure 1 which shows a sample item and some information associated in bullet form from an ecommerce website. The second bullet contains the terms "soups", "casseroles" and "meat" because of which, the item (mushroom) will be present in the recall set for the search queries containing tokens like "soups" and "casseroles" due to full text match, leading to poor search relevancy. Relevant features for this SKU can be thought of as attributes that could be used in a search query to find this product. Thus, "gluten free" and "non-GMO" are considered relevant. Based on the attributes

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR 2019 eCom, July 2019, Paris, France

© 2019 Copyright held by the owner/author(s).




Figure 1: Sample SKU Image and Bullets

present in the sentence and how well it describes the item, we consider third bullet as more relevant to the item than the second bullet from search perspective. Tokens highlighted in Figure 1 with red and green color denote irrelevant and relevant features respectively for the SKU.

Typically, this problem tends to appear in fields like product description and bullets which are often verbose and contain information about the SKU (stock keeping unit, a term used to describe item sold on the site) that is not pertinent to the item. Circuitous descriptions of the product and *Keyword stuffing* are real concern in ecommerce. *Keyword stuffing* refers to the practice of loading product data with keywords that may not be relevant with the item being sold. Figure 2, which is a description of a SKU, illustrates this.

Product descriptions also tend to contain negations. That is, they describe what the product is NOT and what it is not suitable for. These kind of sentences are technically legitimate but poses a challenge for search engines and have the effect of returning misleading or irrelevant results.

A naive solution is to ignore these fields completely for search. While this may improve precision, it would be at the cost of recall, as relevant information might be lost. Such relevancy problems are mitigated by having semantic search using methods like query understanding. However, they require SKUs to have relevant attributes (atomic phrases that provide more information about an item [10]) present in them to match it with user's intent. Thus, attribute extraction from the catalog data is often done in order to enrich SKUs (documents) with relevant attributes.

## Description

Apple Watch Stand, LAMEEKU Desktop Cell Phone Stand For all Android Smartphone, iPhone X 6 6s 7 8 Plus, Samsung, Apple Watch 38mm 42mm, iPad Airpods and Tablet - Black ?WIDE COMPATIBILITY: Universal for apple watch for both 38mm and 42mm models, compatible with ipad, tablet, nintendo switch and all mobile phone (like iPhone, Samsung and so on), although with case. Intimate design, avoids your airpods and apple pencil missing.\*SMART LIFESTYLE: Great desk accessories in the office and house. Perfect 60 angle of view when using facetime, cooking, playing game, watching videos, reading, typing and you can charge your devices at the same time. You would not miss important call, email and message.\*STURDY DESIGN: LAMEEKU stand made of high density wood and 4mm thickness solid aluminum alloy, which is 25 times thicker than normal stand (3mm). Design with low gravity center, lightweight but sturdy. Stand only, charging accessories and cable NOT included.\*RUBBER PROTECTED: Rubber pads and feet protect your device from daily scratches and sliding. The rubber circle on the back of the stand allows you to neatly organize your cables when charging devices or for data transfer to your smart phone.\*WHY SPECIAL: LAMEEKU stand design with aluminum alloy material and nature wood, exquisite workmanship, simple, functional and elegant. 180 days warranty LAMEEKU brand guarantee, as well as our Lifetime Warranty Service.

Features: 1?Simply put you cell phone down and enjoy best using angle 2?Simply put your watch down and charge it automatically 3?Simply put your auroids case on the charging interface and put your Airpods into then charging starts. New Stand VS Old Stand 1. Updated Solid. 2. Wide Compatibility. 3. Anti-Scratch and Anti-Slippery. 4. Multiple Angle Viewing. 5. Excellence on details. The cushy, comfy, cozy stand is designed to keep you and your iPad comfortable. A soft, huggable pillow lets you rest your iPad on your lap or stomach for all those times you'd prefer to skip the desk.LAMEEKUprovides ultimate versatility for your iPad use, all with the comfort of your body in mind.LAMEEKU is designed for you, not just your iPad. Packing List: 1x Metal stand Click theLAMEEKU to find more styles,thanks ?PROP 65 WARNING: This product can expose you to some chemicals, which is known to the State of California to cause For more information, go to www.P65Warnings.ca.gov.?

Figure 2: Sample SKU Description

In this paper we describe a method to rank sentences based on if they are relevant from search perspective, and select top K sentences for search from these fields. Top K ranked set of sentences can lead to better full text match and can also help in extracting attributes for developing the ontology for semantic search [10] as higher ranked sentences would have larger probability of attributes correctly describing the product. In our experiment, we limit K to 3. Thus, given a description of length greater than three, we always pick top three sentences generated by the model as our final summary.

Our contribution in this paper is, we demonstrate how *Extractive Summarization* can be used to rank sentences present in product description and bullets using product title and user queries obtained from click through log. One of the benefits of this method is, cost of obtaining training data is cheap and the model can be run on items that have little or no click data associated with it. We also provide comparison of the two models by measuring precision@k of relevant sentences in the summary.

## 2 RELATED WORK

Summarization is the process of shortening a text document in order to create a summary while retaining major points of the original document. There are two kinds of summarization techniques: *Abstractive* and *Extractive summarization*. *Abstractive summarization* involves using internal semantic representation and natural language generation techniques to create the summary [2] [23], [24]. *Extractive summarization* involves selecting existing subset of words, phrases and sentences in the original text to generate the summary [5], [14], [28].

Recently, a lot of work has been done on *Abstractive Summarization* using attentional encoder-decoder model that was proposed by Sutskever et. al in [25]. In [15], Nallapati et al. modeled abstractive summarization using Attentional Encoder Decoder Recurrent Neural Networks. While in [20], Paulus et. al introduced a new objective function that combined cross entropy loss with rewards from policy gradient reinforcement learning which improved state of the art in abstractive summarization.

*Extractive Summarization* was traditionally done using hand engineered features, such as sentence position, length [21], words present in the sentence, their part of speech tags, frequency etc [18]. However, with the recent success of encoder-decoder model, it is being used in *Extractive Summarization* as well, such as [3] [17] and [16]. In [3], Cheng et al. developed a framework composed of hierarchical document encoder and attention based extractor for extractive summarization. In [17], Narayan et al. used the hierarchical encoder and attention based decoder to leverage side informations like title, image caption etc. and in [16] they introduced a new objective function based on ROUGE and used reinforcement learning to optimize it.

In this paper, we try to rank sentences using summarization techniques for the purpose of improving search relevancy. There hasn't been lot of work done in this area. One of the work that is aligned with our objective is from Ryen et. al [26] published in 2002. They use statistical measures like frequency of query terms present in the sentence to rank them, and recommend user documents from the recall set by presenting them with ranked set of sentences for web search. However, our work focuses on ecommerce setting where we leverage Reinforcement Learning paradigm to rank sentences with the purpose of improving search by affecting recall/precision.

## 3 PROBLEM FORMULATION

Our objective is to rank sentences in product description and bullets from a search perspective. Search perspective means that when we extract attributes from sentences, they are relevant to the item and are likely to be used in a search query for that item. Methods like query understanding can benefit from ranked sentences as they use attributes in SKU to match with the user's intent. Higher ranked sentences are more likely to contain relevant attribute than a lower ranked sentences. Having a set of top ranked sentences would also help in full text match by avoiding queries to match with irrelevant sentences. We use *Extractive Summarization* to achieve this. Our work is based on [16] which treats summarization task as a ranking problem and training is done by optimizing combination of ROUGE metric and cross entropy using reinforcement learning (described in 3.2). ROUGE stands for Recall-Oriented Understudy for Gisting Evaluation. It is a metric to compare automatically generated summary with the reference summary. ROUGE makes use of the count of overlapping units such as N-gram between the two summaries to measure the quality of system generated summary [13]. Here we specifically use F1 score of three ROUGE scores mentioned below:

- ROUGE-1: refers to the overlap of 1-gram between candidate summary and the reference summary (in our case title and queries)
- ROUGE-2: refers to the overlap of bi-gram
- ROUGE-L: measures Longest Common Subsequence based statistics to compute similarity between the two summaries

We use ROUGE because it is well aligned with our objective of finding relevant sentences from SKU description and bullets that is similar to the title and user engagement data (queries). It is the evaluation metric used in most summarization system, and training the

model on a combination of ROUGE and cross entropy is shown to be superior than using just cross entropy [16]. REINFORCE algorithm is shown to improve sequence to sequence based text rewriting systems by optimizing non-differentiable objective function like ROUGE [22] [12], so we use reinforcement learning to optimize our reward function.

We use title and queries obtained from click through log as part of the target summary. Title is one of the key fields in ecommerce catalog provided by the merchant, it captures essential information about the item and queries can be thought of as keywords users think are relevant attributes for the product. The intuition is, having them in the target summary would allow the model to capture important sentences present in the description and bullets. We create two models, one that uses just the title as target summary and the second model that uses top five queries that led to clicks on the item, along with the title as target summary.

Finally, we choose top K sentences as determined by the model as our final summary. Since, ecommerce product description tend to be short and less repetitive, the issue of repetition and diversity in not a concern in our summarization task.

### 3.1 Network Architecture

Figure 3 depicts network architecture of the extractive summarizer. It aims to extract sentences  $\{s_1..s_m\}$  from a document  $D$  composed of sentences  $\{s_1..s_n\}$  where  $n > m$  and labels them 1 or 0 based on if they should be included in the summary or not. It learns to assign a score  $p(y_i|s_i, D, \theta)$  to each sentence which is directly proportional to its relevance within the summary. Here,  $\theta$  denotes the model parameter,  $s_i$  denotes the  $i^{th}$  sentence and  $D$  represents the document. Summary is chosen by selecting the sentences with top  $p(y_i|s_i, D, \theta)$  score. Our network and the objective function is based on the paper [16]. We choose a sequence to sequence network which is composed of three main components: sentence encoder, document encoder and sentence extractor.

These components are described in detail below:

*Sentence encoder* is composed of convolutional encoder which encodes a sentence into a continuous representation and is shown to capture salient features [4], [9], [8]. The encoding is performed using kernel filter  $K$  of width  $h$  over a window of  $h$  words present in the sentence  $s$ . This is applied to each possible window of words in the sentence  $s$  to produce a feature map  $f \in R^{k-h+1}$ , where  $k$  is the length of the sentence. Then max pooling is performed over time on the feature maps and max value is taken corresponding to this particular filter  $K$ . Specifically, we use filter of size 2 and 4.

*Document encoder*: The output of sentence encoder is fed to document encoder. It composes sequence of sentences to obtain a document representation. We use LSTM to achieve this. Given a document  $D$  and sequence of sentence  $(s_1..s_n)$  we feed sentences in reverse order to the model. This approach allows the model to learn that the top sentences are more important and has been

demonstrated in previous work [25], [11], [17].

Finally, *Sentence extractor* sequentially labels each sentence as 1 or 0 depending upon if the sentence is relevant or not. It is implemented using RNN with LSTM cells and a softmax layer. At time  $t_i$ , it makes a binary prediction conditioned on the document representation and previously labelled sentences. This lets it identify locally and globally important sentences. Sentences are then ranked by the score  $p(y_i = 1|s_i, D, \theta)$ . Here  $s_i$  is  $i^{th}$  sentence,  $D$  is the document,  $\theta$  is the model parameter and  $p(y_i = 1|s_i, D, \theta)$  is the probability of sentence  $s_i$  being included in the summary. We learn to rank by training the network in a reinforcement learning framework optimizing ROUGE.

We use a combination of maximum likelihood cross entropy loss and rewards from policy gradient reinforcement learning as objective function to globally optimize ROUGE. This lets the model optimize the evaluation metric directly and makes it better at discriminating sentences i.e it ranks the sentence higher if it appears often in the summary.

### 3.2 Policy Learning

Reinforcement Learning is an area of machine learning where a software agent learns to take actions in an environment to maximize cumulative reward. It differs from supervised learning in the way that labelled input/output pairs need not be provided nor are sub-optimal actions need to be explicitly corrected. Rather, the focus is on the balance between exploration and exploitation. Exploitation is the act of preferring an action that it has tried in the past and was found to be effective, whereas exploration is the act of discovering such actions, i.e. trying out actions that it has not selected before.

We conceptualize the summarization model in a reinforcement learning paradigm. The model can be thought of as an agent interacting with the environment, which consists of documents. The agent reads the document  $D$  and assigns a score to each sentence  $s_i \in D$  using the policy  $p(y_i|s_i, D, \theta)$ . We then rank and get the sampled sentences as the summary. The agent is then given a reward based on how close the generated summary is with the gold standard summary. We use F1 score of ROUGE-1, ROUGE-2, and ROUGE-L as the reward  $r$ . In our case, gold standard summary is the title and user queries. Agent is then updated based on the reward using the REINFORCE algorithm [27]. REINFORCE algorithm minimizes negative expected reward:

$$L(\theta) = -\mathbb{E}_{\hat{y} \sim p_\theta} [r(\hat{y})]$$

Here,  $p_\theta$  stands for  $p(y|D, \theta)$ , where  $\theta$  is the model parameter,  $D$  is the document and  $r$  is the reward.

REINFORCE algorithm is based on the fact that the expected reward function of a non differentiable function can be computed as:

$$\nabla L(\theta) = -\mathbb{E}_{\hat{y} \sim p_\theta} [r(\hat{y}) \nabla \log p(\hat{y}|D, \theta)]$$

Calculating expected gradient in the above expression can be expensive as each document can have very large number of candidate summaries. It can be approximated by taking single sample  $\hat{y}$  from

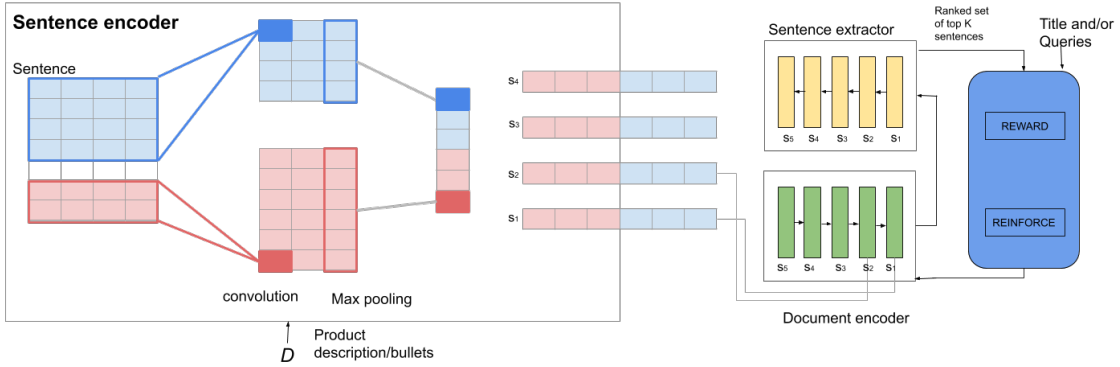


Figure 3: Network architecture

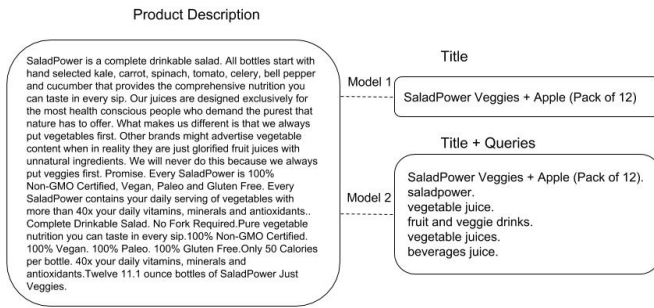


Figure 4: Retrieval by matching query understanding with SKU understanding

$p_\theta$  for each training example in a batch, following which the above expression gets simplified to:

$$\begin{aligned} \nabla L(\theta) &\approx -r(\hat{y}) \nabla \log p(\hat{y}|D, \theta) \\ &\approx -r(\hat{y}) \sum_{i=1}^n \nabla \log p(\hat{y}_i|s_i, D, \theta) \end{aligned}$$

Since the REINFORCE algorithm starts with a random policy, and because our task can involve large number of candidate summaries for the document, training the model can be time consuming. So, we limit the search space  $\hat{y}$  with smaller number of high probability samples  $\hat{Y}$  consisting of top  $k$  extracts. The way we choose these top  $k$  extracts is, we select  $p$  sentences which have highest ROUGE scores on its own and then generate all possible set of combination using these  $p$  sentences with the constraint that maximum length of the extract can be  $m$ . We rank these against the gold summary using F1 score by taking mean of ROUGE-1, ROUGE-2 and ROUGE-L. We choose top  $k$  of these ranked summaries as  $\hat{Y}$ . During training, we sample  $\hat{y}$  from  $\hat{Y}$  instead of  $p(\hat{y}|\theta, D)$ .

### 3.3 Input Data for model

We create two summarization models, one with title as its target summary (Model 1) and the other with title plus top five queries for which the product was clicked as the target summary (Model 2). Title and each query are treated as independent sentences when




Figure 5: Precision @ k for the baseline

generating the reference summary. For input, we use product descriptions and bullets for both the models.

We preprocess the title, description and queries before passing them to the model. Preprocessing step consists of sentence segmentation, tokenization, conversion of tokens into vocabulary id, truncation and padding to a fixed length. We use SKUs from grocery category of our catalog to evaluate the models. For *Model 1* we used all the SKUs from the grocery category and for *Model 2* we used a subset of SKUs from the category which had engagement above a certain threshold. Though *Model 2* had fewer training data, it was richer since it had queries (top 5) associated with each SKU as part of the summary. One advantage of both methods is, it requires almost no manual effort to get the training data, thus is very cheap. Figure 4 describes how the two models are set up for training.

Since our objective is to have better full text match or attributes from the ranked set of sentences, each sentence can be independent of each other. This insight is well aligned with the framework of reinforcement learning based extractive summarization that optimizes ROUGE.

## 4 BASELINE MODEL

*Tfidf* is one of the commonly used frequency driven approaches for weighting terms to measure importance of a sentence for extractive summarization [1], [19]. It measures the importance of words and identifies very common words in the documents by giving low weights to words appearing in most documents. The weight of each word is computed by the formula:

$$tfidf(t, d, D) = tf(t, d) \cdot idf(t, D)$$

$$idf(t, D) = \frac{N}{|\{d \in D : t \in d\}|}$$

Here,  $tf(t, d)$  is the count of the term  $t$  in the document  $d$ .  $idf(t, D)$  is the inverse document frequency.  $N$  is the total number of documents in the corpus.  $|\{d \in D : t \in d\}|$  is the number of documents where the term  $t$  appears. If the term is not present in the corpus, it will lead to division by zero. To avoid this, it is a common practice to adjust the denominator to  $1 + |\{d \in D : t \in d\}|$ .

For baseline, we use *tfidf* based model. Our baseline consists of three approaches that utilizes *tfidf* to score the sentences to select top  $K$ . For the first approach, we sum up (unweighted) *tfidf* score of the words to measure importance of a sentence and then select top  $K$  as the summary. Here,  $tf$  is computed at the sentence level and  $idf$  is across all the SKUs (documents).

For the second approach (weighted), we weigh the *tfidf* score of tokens in the description that also appear in the title by multiplying it with a factor of  $w_i$ . The optimal weight  $w_i$  was found by using grid search method. In our case, it was found to be 2.

For the third approach (filtered), we sum up the *tfidf* score of only those tokens in description that appear in the title.

Figure 5 shows precision@k for the three models. As we can see from the graph, the weighted approach has highest precision@k, this shows that the words present in title does indicate which sentences are of relatively higher importance. However, it is also not a right strategy to exclude all the other words, as demonstrated by the higher precision@k of unweighted model over filtered model. Thus, in summary, boosting words present in title while also retaining other words for the computation of *tfidf* score of a sentence seems to yield best result among all the baseline approaches.

## 5 EVALUATION

Our purpose of ranking is to find sentences that are relevant to the product and contain attributes of the product that customers might use in their search queries. This will improve results of full text match as well as query understanding, since it depends on matching user's intent with attributes extracted from the SKU. To analyze this, we reviewed 100 SKUs randomly sampled from the grocery category and manually labeled the sentences based on whether they were relevant or not. We evaluated the model using precision@k, with  $k$  as 1, 2 and 3.

Based on the evaluation of the three *tfidf* based models as described in the section 4, we chose *weighted* Model (the second approach)




Figure 6: Precision @ k for Model 1(Title only), Model 2 (Title and queries) and the baseline

as our baseline, as it has the best performance.

Figure 6 shows precision@k for the two sequence to sequence based model and the baseline. Blue line indicates the model that was trained using just the title as target summary (Model 1), orange line indicates the model that was trained using title and top five queries that led to clicks on the SKU (Model 2) while, gray line is the precision@k for the baseline. We found that both Model 1 and Model 2 outperform the baseline. Model 2 was better by 3.125% and 12.08% over Model 1 for precision@2 and precision@3 respectively. We believe the reason for Model 2 to outperform Model 1 is that queries provide additional context regarding which sentences are important and captures key information of the product, which is key to summarization.

This demonstrates that words present in title capture key information of the product being sold. Title is provided by the merchant, so it provides merchant's point of view regarding what aspect of the product is important. Whereas, words present in user queries indicate the attributes of product that the user cares about. So combining these two sources of information is a good way to infer relevant sentences of description from a search perspective. Also, since not all SKUs (documents) have user clicks or may have comparatively less engagement data associated with it, creating a model leveraging title and click through log to find relevant sentences provides a way to generalize it to SKUs (documents) that have little or no engagement data.

We provide one instance from our evaluation set as an example. Figure 7 shows a sample product description that is fed to the model. Figures 8 and 9 show output of Model 1 and Model 2 respectively.

Sentences that have keyword stuffing tend to be grammatically incorrect, structurally dissimilar to the title and generally longer. Thus, the intuition is that summarization models described above



Food For Life's unique Wheat and Gluten Free Multi Seed English Muffins are specifically developed to be moist, flavorful and easy to use. We believe in keeping things pure and simple. These Gluten Free English muffins are made with your health in mind and contain only natural ingredients. Adding whole grains and seeds to your diet makes sense. Plus they taste great toasted! You can prepare these Gluten Free Multi-Seed English Muffins the same as any English muffin. Simply gently separate, toast and enjoy hot with your favorite topping. Add fruit spread or nut butter (as desired). With just one bite, you'll know they're a food for life. This product arrives frozen to preserve the best flavor and quality. Gluten Free. Vegan. All Natural.

**Figure 7: Input to the model: product description**

1. Food For Life's unique Wheat and Gluten Free Multi Seed English Muffins are specifically developed to be moist, flavorful and easy to use.
2. We believe in keeping things pure and simple.
3. These Gluten Free English muffins are made with your health in mind and contain only natural ingredients.

**Figure 8: Model 1's output (title)**

1. Food For Life's unique Wheat and Gluten Free Multi Seed English Muffins are specifically developed to be moist, flavorful and easy to use.
2. Add fruit spread or nut butter -LRB- as desired -RRB-.
3. Vegan.

**Figure 9: Model 2's output (title + query)**

would rank such sentences lower.

## 6 CONCLUSION AND FUTURE WORK

We implemented a framework to rank sentences from product description & bullets based on *Extractive Summarization* that uses reinforcement learning to optimize ROUGE and maximum likelihood cross entropy, thus enabling the model to learn rank the sentences. We compare two models, one that uses just the title and the other that uses queries from click through log along with the title. We show that these two models have higher precision in finding relevant sentences than the baseline which is a tf-idf based method to select top sentences. Typically, in search engines, such fields (product descriptions, bullets etc.) are either ignored or given a very low weight compared to fields like product title. Using this framework that ranks the sentences, we can assign a higher weight to ranked set of sentences. In addition, top N sentences from ranked set can also be used to extract attributes and help build the ontology.

Our future plan involves, 1) measuring the precision with two separate models, one for description and one for bullets, as they tend to have different grammatical structure 2) investigate the effect of query length on the ranking of sentences 3) have an algorithmic method to decide on the cut off (Top N) for selecting top sentences from each SKU. This is because, as length of the content in each SKU varies, number of relevant sentences could be different.

## REFERENCES

- [1] Mehdi Allahyari, Seyedamin Pouriyeh, Mehdi Assefi, Saied Safaei, Elizabeth D Trippe, Juan B Gutierrez, and Krysz Kochut. 2017. Text summarization techniques: a brief survey. *arXiv preprint arXiv:1707.02268* (2017).
- [2] Qian Chen, Xiaodan Zhu, Zhenhua Ling, Si Wei, and Hui Jiang. 2016. Distraction-based neural networks for document summarization. *arXiv preprint arXiv:1610.08462* (2016).
- [3] Jianpeng Cheng and Mirella Lapata. 2016. Neural summarization by extracting sentences and words. *arXiv preprint arXiv:1603.07252* (2016).
- [4] Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of machine learning research* 12, Aug (2011), 2493–2537.
- [5] Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of artificial intelligence research* 22 (2004), 457–479.
- [6] Clinton Gormley and Zachary Tong. 2015. *Elasticsearch: The Definitive Guide: A Distributed Real-Time Search and Analytics Engine*. " O'Reilly Media, Inc".
- [7] Trey Grainger, Timothy Potter, and Yonik Seeley. 2014. *Solr in action*. Manning Cherry Hill.
- [8] Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. *arXiv preprint arXiv:1404.2188* (2014).
- [9] Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882* (2014).
- [10] Aliasgar Kutiyawala, Prateek Verma, and Zheng Yan. 2018. Towards a simplified ontology for better e-commerce search. *CoRR abs/1807.02039* (2018). arXiv:1807.02039 <http://arxiv.org/abs/1807.02039>
- [11] Jiwei Li, Minh-Thang Luong, and Dan Jurafsky. 2015. A hierarchical neural autoencoder for paragraphs and documents. *arXiv preprint arXiv:1506.01057* (2015).
- [12] Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016. Deep reinforcement learning for dialogue generation. *arXiv preprint arXiv:1606.01541* (2016).
- [13] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out* (2004).
- [14] Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [15] Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023* (2016).
- [16] Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Ranking sentences for extractive summarization with reinforcement learning. *arXiv preprint arXiv:1802.08636* (2018).
- [17] Shashi Narayan, Nikos Papasrantopoulos, Shay B Cohen, and Mirella Lapata. 2017. Neural extractive summarization with side information. *arXiv preprint arXiv:1704.04530* (2017).
- [18] 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*. ACM, 573–580.
- [19] Joel Larocca Neto, Alexandre D Santos, Celso AA Kaestner, Neto Alexandre, D Santos, et al. 2000. Document clustering and text summarization. (2000).
- [20] Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304* (2017).
- [21] Dragomir R Radev, Timothy Allison, Sasha Blair-Goldensohn, John Blitzer, Arda Celebi, Stanko Dimitrov, Elliott Drabek, Ali Hakim, Wai Lam, Danyu Liu, et al. 2004. MEAD—a platform for multidocument multilingual text summarization. (2004).
- [22] Marc'Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732* (2015).
- [23] Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. *arXiv preprint arXiv:1509.00685* (2015).
- [24] Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368* (2017).
- [25] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*. 3104–3112.
- [26] Ryen W White, Ian Ruthven, and Joemon M Jose. 2002. Finding relevant documents using top ranking sentences: an evaluation of two alternative schemes. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 57–64.
- [27] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8, 3-4 (1992), 229–256.
- [28] Michihiro Yasunaga, Rui Zhang, Kshitij Meelu, Ayush Pareek, Krishnan Srinivasan, and Dragomir Radev. 2017. Graph-based neural multi-document summarization. *arXiv preprint arXiv:1706.06681* (2017).